

The Alive Particle Filter

BY AJAY JASRA¹, ANTHONY LEE², CHRIS YAU³ & XIAOLE ZHANG³

¹Department of Statistics & Applied Probability, National University of Singapore, Singapore, 117546, SG.

E-Mail: *staja@nus.edu.sg*

²Department of Statistics, University of Warwick, Coventry, CV4 7AL, UK.

E-Mail: *anthony.lee@warwick.ac.uk*

³Department of Mathematics, Imperial College London, London, SW7 2AZ, UK.

E-Mail: *c.yau@ic.ac.uk, x.zhang11@ic.ac.uk*

Abstract

In the following article we develop a particle filter for approximating Feynman-Kac models with indicator potentials. Examples of such models include approximate Bayesian computation (ABC) posteriors associated with hidden Markov models (HMMs) or rare-event problems. Such models require the use of advanced particle filter or Markov chain Monte Carlo (MCMC) algorithms e.g. Jasra et al. (2012), to perform estimation. One of the drawbacks of existing particle filters, is that they may ‘collapse’, in that the algorithm may terminate early, due to the indicator potentials. In this article, using a special case of the locally adaptive particle filter in Lee et al. (2013), which is closely related to Le Gland & Oudjane (2004), we use an algorithm which can deal with this latter problem, whilst introducing a random cost per-time step. This algorithm is investigated from a theoretical perspective and several results are given which help to validate the algorithms and to provide guidelines for their implementation. In addition, we show how this algorithm can be used within MCMC, using particle MCMC (Andrieu et al. 2010). Numerical examples are presented for ABC approximations of HMMs.

Key Words: Particle Filters, Markov Chain Monte Carlo, Feynman-Kac Formulae.

1 Introduction

Let $\{(\mathcal{E}_n, \mathcal{E}_n)\}_{n \geq 1}$ be a sequence of measurable spaces, $\{G_n(x) = \mathbb{I}_{B_n}(x)\}_{n \geq 1}$, $(x, B_n) \in E_n \times \mathcal{E}_n$, $B_n \subset E_n$, be a sequence of indicator potentials and $\{M_n : E_{n-1} \times \mathcal{E}_n \rightarrow [0, 1]\}_{n \geq 1}$, with $x_0 \in E_0$ a fixed point, be a sequence of Markov kernels. Then for the collection of bounded and measurable functions $\varphi \in \mathcal{B}_b(E_n)$ the n -time Feynman-Kac marginal is:

$$\eta_n(\varphi) := \frac{\gamma_n(\varphi)}{\gamma_n(1)}, \quad n \geq 1$$

assuming that $\gamma_n(\varphi) = \mathbb{E}_{x_0}[\prod_{p=1}^{n-1} G_p(X_p)]$ is well-defined, where $\mathbb{E}_{x_0}[\cdot]$ is the expectation w.r.t. the law of an inhomogeneous Markov chain with transition kernels $\{M_n\}_{n \geq 1}$. Such models appear routinely in the statistics and applied probability literature including:

- ABC approximations (as in, e.g., Del Moral et al. (2012))
- ABC approximations of HMMs (Dean et al. 2010; Jasra et al. 2012)
- Rare-Events problems (as in, e.g., Cérou et al. (2012))

In order to perform estimation for such models, one often has to resort to numerical methods such as particle filters or MCMC; see the aforementioned references.

The basic particle filter, at time n and given a collection of samples $N \geq 1$ with non-zero potential on E_{n-1}^N , will generate samples on E_n using the Markov kernels $\{M_n\}_{n \geq 1}$ and then sample with replacement amongst $\{x_n^i\}_{1 \leq i \leq N}$ according to the normalized weights $G_n(x_n^i) / \sum_{j=1}^N G_n(x_n^j)$. The key issue with this basic particle filter is that, at any given time, there is no guarantee that any sample x_n^i lies in B_n , and in some challenging scenarios, the algorithm can ‘die-out’ (or collapse), that is, that all of the samples have zero potentials. From an inference perspective, this is clearly an undesirable property and can lead to some poor performances; for example, many particle filters display time-uniform convergence properties, which has yet to be shown for this class of particle filters. For some classes of examples, e.g. Cérou et al. (2012) or Del Moral et al. (2012), there are some adaptive techniques which can reduce the possibility of the algorithm collapsing, but these are not always guaranteed to work in practice. In this article we develop a particle filter. This algorithm uses the same sampling mechanism, but the samples are generated until there is a prespecified number that are alive. This removes the possibility that the algorithm can collapse, but introduces a random cost per time-step. The algorithm turns out to be an important special case of the work in Lee et al. (2013) and is closely related to Le Gland & Oudjane (2004).

The particle filter is analyzed from a theoretical perspective. In particular, under assumptions, we establish the following results:

1. Time uniform \mathbb{L}_p bounds for the particle filter estimates of $\eta_n(\varphi)$
2. A central limit theorem (CLT) for suitably normalized and centered particle filter estimates of $\eta_n(\varphi)$
3. An unbiased property of the particle filter estimates of $\gamma_n(\varphi)$
4. The relative variance of the particle filter estimates of $\gamma_n(\varphi)$, assuming $N = \mathcal{O}(n)$, is shown to grow linearly in n .

Whilst all of these results are classical in the literature on particle filters (C  rou et al. 2011; Del Moral 2004), the proof in this new context requires some modifications. In the main, these technical adjustments are associated to \mathbb{L}_p -bounds and CLTs for sums of random variables with a random number of summands (in the context of 1.-2.). The technical results in 1.-2. not only verify the correctness of the new algorithm, but suggest a substantial improvement over the standard particle filter, at the cost of increased computational time. The results in 3.-4. are of particular interest when using the new particle filter within MCMC methodology (a particle MCMC (PMCMC) algorithm, (Andrieu et al. 2010)). There are variety of applications of such PMCMC algorithms, for example, when performing static parameter estimation for ABC approximations of HMMs. The results in 3.-4. not only allow one to construct new PMCMC algorithms, but also provide theoretical guidelines for their implementation. It is remarked that some of these results can also be found in Le Gland & Oudjane (2006) (with regards to 1. 2.), except for a *different estimate*; this is a critical difference between the work in this article and that in Le Gland & Oudjane (2006). In particular, as mentioned in result 3. our estimates of $\gamma_n(\varphi)$, and in particular $\gamma_n(1)$ are unbiased and this allows one to develop principled MCMC methodology. We also note that, in Le Gland & Oudjane (2006) the authors do not give a time-uniform bound.

The structure of this article is as follows. In Section 2 we provide a motivating example, ABC approximations of HMMs, for the construction of the particle filter, as well as the new particle filter itself. In Section 3 our theoretical results are provided along with some interpretation of their meaning. In Section 4 we implement the new particle filter for the motivating example and then develop a basic PMCMC algorithm using the guidelines in Section 3 for static parameter estimation associated to ABC approximations of HMMs. In Section 5 the article is concluded, with some discussion of future work. The appendix contains technical results for the theory in Section 3 and is split into three sections.

2 Motivating Example and Algorithm

2.1 Motivating Example

We are given a HMM with observations $\{Y_n\}_{n \geq 1}$, $Y_n \in \mathcal{Y} \subseteq \mathbb{R}^{d_y}$, hidden states $\{Z_n\}_{n \geq 0}$, $Z_n \in \mathcal{X} \subseteq \mathbb{R}^{d_x}$, Z_0 given. We assume:

$$\mathbb{P}(Y_n \in A | \{Z_n\}_{n \geq 0}) = \int_A g_\theta(y|z_n) dy \quad n \geq 1$$

and

$$\mathbb{P}(Z_n \in A | \{Z_n\}_{n \geq 0}) = \int_A f_\theta(z|z_{n-1}) dy \quad n \geq 1$$

with $\theta \in \Theta$ a static parameter and dy Lebesgue measure.

We assume $g_\theta(y|x_n)$ is unknown (even up to an unbiased estimate), but one can sample from the associated distribution. In this scenario, one cannot apply a standard particle filter (or many other numerical approximation schemes). Dean et al. (2010) and Jasra et al. (2012) introduce the following ABC approximation of the joint smoothing density, for $\epsilon > 0$:

$$\pi_\theta(z_{1:n}|y_{1:n}) = \frac{\prod_{k=1}^n g_\theta^\epsilon(y_k|z_k) f_\theta(z_k|x_{k-1})}{\int_{\mathcal{X}^n} \prod_{k=1}^n g_\theta^\epsilon(y_k|z_k) f_\theta(z_k|z_{k-1}) dz_{1:n}} \quad (1)$$

where

$$g_\theta^\epsilon(y_k|x_k) = \frac{\int_{B_\epsilon(y_k)} g_\theta(u|x_k) du}{\int_{B_\epsilon(y_n)} du}$$

and $B_\epsilon(y_k)$ is the open ball centered at y_k with radius ϵ .

We let θ be fixed and omit it from our notations; it is reintroduced later on. We introduce a Feynman-Kac representation of the ABC approximation described above. Let $\mathbf{E}_n = \mathbf{E} \times \mathcal{Y}$ and define $G_n : \mathbf{E} \rightarrow \{0, 1\}$:

$$G_n(x) = \mathbb{I}_{\mathcal{X} \times B_\epsilon(y_n)}(x).$$

Now introduce Markov kernels $\{M_n\}_{n \geq 1}$, $M_n : \mathbf{E} \times \mathbb{B}(\mathbf{X} \times \mathbf{Y}) \rightarrow [0, 1]$ ($\mathbb{B}(\cdot)$ are the Borel sets), with

$$M_n(x, dx') = f(z'|z)g(u'|z')du'dz'$$

with $x = (z, u)$. Then the ABC predictor is for $n \geq 1$:

$$\eta_n(\varphi) := \frac{\gamma_n(\varphi)}{\gamma_n(1)}, \quad (2)$$

where $\varphi \in \mathcal{B}_b(\mathbf{E})$ and

$$\gamma_n(\varphi) = \mathbb{E}_{x_0} \left[\prod_{p=1}^{n-1} G_p(X_p) \varphi(X_n) \right] = \int_{\mathbf{E}^n} \prod_{p=1}^{n-1} G_p(x_p) M_p(x_{p-1}, dx_p) M_n(x_{n-1}, dx_n) \varphi(x_n). \quad (3)$$

This provides a concrete example of the Feynman-Kac model in Section 1. In light of (2), we henceforth refer to $\gamma_n(1)$ as the *normalizing constant*. This quantity is of fundamental importance in a wide variety of statistical applications, notably in static parameter estimation, as it is equivalent to the marginal likelihood of the observed data Y_1, \dots, Y_{n-1} in contexts such as the ABC approximation presented above, as can be determined from (3).

2.2 Old Filter

Now define, for $n \geq 2$:

$$\Phi_n(\eta_{n-1})(\varphi) = \frac{\eta_{n-1}(G_{n-1}M_n(\varphi))}{\eta_{n-1}(G_{n-1})}.$$

The standard particle filter works by sampling x_1^1, \dots, x_1^N i.i.d from $M_1(x_0, \cdot)$ and setting

$$\eta_n^N(\varphi) = \frac{1}{N} \sum_{i=1}^N \varphi(x_n^i) \quad n \geq 1$$

at times $n \geq 2$ sampling x_n^1, \dots, x_n^N from $\Phi_n(\eta_{n-1}^N)(\cdot)$, *assuming that the system has not died out*.

2.3 Alive Filter

We now discuss an idea which will prevent the particle filter from dying out; see also Lee et al. (2013) and Le Gland & Oudjane (2006). Throughout we assume that $M_n(x, \mathbf{B}_n)$ is not known for each x, n ; if this is known, then one can develop alternative algorithms. At time 1, we sample $x_1^1, \dots, x_1^{T_1}$ i.i.d. from $M_1(x_0, \cdot)$, where

$$T_1 = \inf\{n \geq N : \sum_{i=1}^n G_1(x_1^i) \geq N\}.$$

Then, define

$$\eta_1^{T_1}(\varphi) = \frac{1}{T_1 - 1} \sum_{i=1}^{T_1-1} \varphi(x_1^i).$$

Now, at time 2 sample $x_2^1, \dots, x_2^{T_2}$, conditionally i.i.d. from $\Phi_2(\eta_1^{T_1})(\cdot)$, where

$$T_2 = \inf\{n \geq N : \sum_{i=1}^n G_2(x_2^i) \geq N\}.$$

This is continued until needed (i.e. with an obvious definition of T_3, T_4 etc). The idea here is that, at every time step, we retain $N - 1$ particles with non-zero weight, so that the algorithm never dies out, but with the additional issue that the computational cost per time-step is a random variable. The procedure is described in Algorithm 1. We note that the approach in Le Gland & Oudjane (2004) retains N alive particles, i.e. it differs only in step 2(a) of Algorithm 1 by sampling instead a_{p-1}^j uniformly on $\{k \in \{1, \dots, T_{p-1}\} : G_{p-1}(x_{p-1}^k) = 1\}$. This seemingly innocuous difference is, however, crucial to the unbiasedness results we develop in the sequel.

Algorithm 1 Alive Particle Filter

1. At time 1. For $j = 1, 2, \dots$ until $j =: T_1$ is reached such that $G_1(x_1^j) = 1$ and $\sum_{i=1}^j G_1(x_1^i) = N$:
 - Sample x_1^j from $M_1(x_0, \cdot)$.
 2. At time $1 < p \leq n$. For $j = 1, 2, \dots$ until $j =: T_p$ is reached such that $G_p(x_p^j) = 1$ and $\sum_{i=1}^j G_p(x_p^i) = N$:
 - (a) Sample a_{p-1}^j uniformly from $\{k \in \{1, \dots, T_{p-1} - 1\} : G_{p-1}(x_{p-1}^k) = 1\}$.
 - (b) Sample x_p^j from $M_p(x_{p-1}^{a_{p-1}^j}, \cdot)$.
-

2.3.1 Some Remarks

We remark that one can show (Del Moral, 2004) that for $n \geq 2$, the normalizing constant is given by

$$\gamma_n(1) = \prod_{p=1}^{n-1} \eta_p(G_p).$$

Thus, a natural estimate of the normalizing constant is

$$\gamma_n^{T_n}(1) = \prod_{p=1}^{n-1} \eta_p^{T_p}(G_p) = \prod_{p=1}^{n-1} \frac{N-1}{T_p-1}.$$

We note that the estimates of η_n and γ_n are different from those considered in Le Gland & Oudjane (2006). This is a critical point as in Proposition 3.1 we show that this estimate of the normalizing constant is unbiased which is crucial for using this idea inside MCMC algorithms. In this direction, one uses the particle filter to help propose values and there is an accept/reject step; we discuss this approach in Section 4.2. Again, it is clearly undesirable in an MCMC proposal, if the particle filter will collapse and so, our approach will prove to be very useful in this context.

Other than the fact that this filter will not die out, in the context of our motivating example, there is also a natural use of this idea. This is because, one can envisage the arrival of an outlier or unusual data; in such scenarios, the alive particle filter will assign (most likely) more computational effort for dealing with this issue, which is not something that the standard filter is designed to do.

A final remark is as follows; in our example $\mathbf{B}_n = \mathbf{X} \times B_\epsilon(y_n)$ and so, as assumed in this article in general, $M_n(x, \mathbf{B}_n)$ is not known for each x, n . This removes the possibility of changing measure to \mathbb{Q} (in the formula for $\gamma_n(\cdot)$), (with finite dimensional marginal \mathbb{Q}_n)

$$\mathbb{Q}_n(d(x_1, \dots, x_n)) \prod_{p=1}^n \frac{M_p(x_{p-1}, dx_p) \mathbb{I}_{\mathbf{B}_p}(x_p)}{M_p(x_{p-1}, \mathbf{B}_p)}$$

call the Markov kernels in the product \hat{M}_p . This is because the new potential at time n is exactly: $M_n(x, \mathbf{B}_n)$. However, one can simulate from \hat{M}_n and use an unbiased estimate of $M_n(x, \mathbf{B}_n)$ for each particle. That is, we obtain samples $z^{(1)}, z^{(2)}$ from $\hat{M}_p(x_{p-1}^i, \cdot)$ using R samples (total) from $M_p(x_{p-1}^i, \cdot)$ and then we set $x_p^i = z^{(1)}$ (say) with associated weight $1/(R-1)$. This particular procedure would then have a fixed number of particles with no possibility of collapsing. Other than the algorithm being convoluted, some particles x_{p-1}^i could be such that $\mathbb{E}[R]$ is prohibitively large, even though $\mathbb{E}[T_p]$ is not very large, which provides a reasonable argument against such a scheme.

3 Theoretical Results

We will now present some theoretical results for the particle filter in Section 2.3. This Section could be skipped with little loss of continuity in the article; although we do provide numerical simulations to verify the behaviour that is predicted by the forthcoming theoretical results.

3.1 Assumptions and Notations

Define the following sequence of Markov kernels, for $n \geq 1$:

$$\hat{M}_n(x, dy) = \frac{M_n(x, dy)G_n(y)}{M_n(G_n)(x)}.$$

We will make use of the following assumptions:

- (M_1) : For each there exist a $\delta \in (0, 1)$ such that for each $n \geq 1$, $(x, y) \in \mathbb{E}_n^2$

$$M_n(x, \cdot) \geq \delta M_n(y, \cdot).$$

In addition there exist a $0 < c < 1$ such that for each $n \geq 1$, $x \in \mathbb{E}_{n-1}$, $M_n(\mathbb{B}_n)(x) \wedge M_n(\mathbb{B}_n^c)(x) \geq c$.

- (\hat{G}) : For each $n \geq 0$

$$\sup_{(x,y) \in \mathbb{B}_n^2} \frac{M_{n+1}(G_n)(x)}{M_{n+1}(G_n)(y)} = \delta_n < \infty.$$

- (\hat{M}_m) : There exist $m \geq 1$ and $\beta_p^{(m)} \in [1, \infty)$ such that for any $p \geq 1$ and $(x, y) \in \mathbb{B}_p$

$$\hat{M}_{p,p+m}(x, dz) \leq \hat{\beta}_p^{(m)} \hat{M}_{p,p+m}(x, dz) \quad \hat{M}_{p,p+m} = \hat{M}_{p+1} \hat{M}_{p+2} \dots \hat{M}_{p+m}.$$

The final two conditions are (\hat{H}_m) in Cérou et al. (2011); we also use the notation $\hat{\delta}_p^{(m)} = \prod_{q=p}^{p+m-1} \hat{\delta}_q$. These assumptions are exceptionally strong, but we remark that for the scenario of interest, weaker conditions have not been used in the literature. Note that in addition, in the context of ABC, the assumptions are essentially qualitative as verifying them is very difficult (even on compact state-spaces) as the likelihood density is typically intractable. However, we still expect the phenomena reported in the below results to hold in some practical situations. We again remark that our results are relevant for scenarios other than ABC.

In order to understand some of the subsequent results, we introduce some notations. For a probability measure on \mathbb{E} (denoted $\mathcal{P}(\mathbb{E})$) $\mu \in \mathcal{P}(\mathbb{E})$ and bounded measurable real-valued function (denoted $\mathcal{B}_b(\mathbb{E})$) $\varphi \in \mathcal{B}_b(\mathbb{E})$, we write $\mu(\varphi) := \int_{\mathbb{E}} \varphi(x) \mu(dx)$. For $\varphi \in \mathcal{B}_b(\mathbb{E})$, $\|\varphi\|_{\infty} := \sup_{x \in \mathbb{E}} |\varphi(x)|$. For $\varphi \in \mathcal{B}_b(\mathbb{E})$, $\text{Osc}(\varphi) = \sup_{(x,y) \in \mathbb{E}^2} |\varphi(x) - \varphi(y)|$. For $\mu, \nu \in \mathcal{P}(\mathbb{E})$, $\|\mu - \nu\|_{tv}$ denotes the total variation distance. For a non-negative operator on $\mathcal{B}_b(\mathbb{E})$, $R(x, \cdot)$, and $\varphi \in \mathcal{B}_b(\mathbb{E})$, $R(\varphi)(x) = \int_{\mathbb{E}} \varphi(y) R(x, dy)$. Iterates of R are written $R^n(x_0, dx_n) = \int R(x_0, dx_1) \times \dots \times R(x_{n-1}, dx_n)$. We will use the semi-group $Q_n(x, dy) = G_{n-1}(x) M_n(x, dy)$ with $n \geq 2$, with the convention for $p < n$, $Q_{p,n}(\varphi)(x_p) = \int Q_{p+1}(x_p, dx_{p+1}) \dots Q_n(x_{n-1}, dx_n) \varphi(x_n)$, where $\varphi \in \mathcal{B}_b(\mathbb{E})$; when $p = n$, $Q_{n,n}$ is the identity operator. We also adopt the notation for $\mu \in \mathcal{P}(\mathbb{E})$ $\Phi_{p,n}(\mu)(\varphi) = \Phi_n \circ \dots \circ \Phi_{p+1}(\mu)(\varphi)$, $\varphi \in \mathcal{B}_b(\mathbb{E})$, $p < n$; when $p = n$, $\Phi_{n,n}(\mu)$ is the identity operator. \mathbb{E} denotes expectation w.r.t. the stochastic process which generates the algorithm, with corresponding probability \mathbb{P} . It is assumed that $\prod_{\emptyset} = 1$. Note the important formula $\gamma_n(\varphi) = [\prod_{q=1}^{n-1} \eta_q(G_q)] \eta_n(\varphi)$, $\varphi \in \mathcal{B}_b(\mathbb{E})$. $\mathcal{N}(\mu, \sigma^2)$ denotes the normal distribution with mean μ and variance σ^2 . $\mathcal{Geo}(p)$ denotes a geometric random variable (with support $\{1, 2, \dots\}$) with success probability p .

3.2 Predictor

In this Section, we consider the long-time behaviour of approximation of the prediction filter

$$\eta_n^{T_n}(\varphi) = \frac{1}{T_n - 1} \sum_{i=1}^{T_n-1} \varphi(z_n^i).$$

In particular, the study of this latter behaviour w.r.t. the algorithm in Section 2.2, is difficult due to the fact that the algorithm can collapse. For example, in a slightly different context, it is shown in Del Moral & Doucet (2004) that the algorithm which can die out has an upper-bound on the \mathbb{L}_p -error which increases with n (and under strong hypotheses as in this article). In general, we do not know of any time-uniform result for algorithms which can die out. Below, we restrict $p \in [1, 4]$ as this is all that is needed for a strong law of large numbers. The additional technical results associated to Theorem 3.1 can be found in Appendix A.

Theorem 3.1. *Assume (M_1) . Then for any $p \in [1, 4]$ there exist a $C_p < \infty$ such that for any $n \geq 1$, $N \geq 2$, $\varphi \in \mathcal{B}_b(\mathbb{E}_n)$:*

$$\mathbb{E}[|\eta_n^{T_n}(\varphi) - \eta_n(\varphi)|^p]^{1/p} \leq \frac{C_p \|\varphi\|_{\infty}}{\sqrt{N-1}}.$$

Proof. Throughout C_p is a finite positive constant (that does not depend upon n) whose value may change from line to line. The proof follows that of Theorem 7.4.4 of Del Moral (2004). We have, using eq. 7.24 of Del Moral (2004)

$$\mathbb{E}[|\eta_n^{T_n}(\varphi) - \eta_n(\varphi)|^p]^{1/p} \leq \sum_{q=1}^n \mathbb{E}[|\Phi_{q,n}(\eta_q^{T_q}) - \Phi_{q,n}(\Phi_q(\eta_{q-1}^{T_{q-1}}))|(\varphi)|^p]^{1/p}.$$

WLOG we suppose $\text{Osc}(\varphi) \leq 1$. Now for $x \in \mathcal{B}_q$ we define the Markov kernel $P_{q,n}(x, \cdot) := Q_{q,n}(x, \cdot)/Q_{q,n}(1)(x)$ with associated Dobrushin coefficient $\beta(P_{q,n}) = \sup_{(x,y) \in \mathcal{B}_q^2} \|P_{q,n}(x, \cdot) - P_{q,n}(y, \cdot)\|_{tv}$ and also set $r_{q,n} = \sup_{(x,y) \in \mathcal{B}_q^2} Q_{q,n}(1)(x)/Q_{q,n}(1)(y)$. Then following the calculations of pp.245–246 of Del Moral (2004), we have

$$\mathbb{E}[|\eta_n^{T_n}(\varphi) - \eta_n(\varphi)|^p]^{1/p} \leq \sum_{q=1}^n r_{q,n} \beta(P_{q,n}) \mathbb{E}[|\eta_q^{T_q} - \Phi_q(\eta_{q-1}^{T_{q-1}})|(\bar{Q}_{q,n}^N(\varphi))|^p]^{1/p}.$$

where $\bar{Q}_{q,n}^N(\varphi)$ is defined in pp. 246 of Del Moral (2004) and note that $\|\bar{Q}_{q,n}^N(\varphi)\|_\infty \leq 1$. Application of Corollary A.1 gives:

$$\mathbb{E}[|\eta_n^{T_n}(\varphi) - \eta_n(\varphi)|^p]^{1/p} \leq \frac{C_p}{\sqrt{N-1}} \sum_{q=1}^n r_{q,n} \beta(P_{q,n}).$$

The sum on the R.H.S. can be bounded uniformly in n by using standard arguments in Cérou et al. (2011) or Del Moral (2004) and are hence omitted. This concludes the proof. \square

3.3 Central Limit Theorem

In this Section we consider the asymptotic properties of a suitably normalized and centered estimate of the predictor; a central limit theorem. We note that such a result is not a direct corollary of existing CLTs for particle filters in the literature (e.g. Del Moral (2004)). The additional technical results associated to Theorem 3.2 can be found in Appendix B.

Here we write:

$$\eta_n^{N-1}(\varphi) = \frac{1}{N-1} \sum_{i=1}^N \varphi(X_n^i)$$

the empirical measure of the first $N-1$ sampled particles at time n . The convergence in probability (written $\rightarrow_{\mathbb{P}}$) weak convergence (written \Rightarrow) results are as $N \rightarrow \infty$.

Theorem 3.2. *Assume (M_1) . Then for any $n \geq 1$, $\varphi \in \mathcal{B}_b(\mathcal{E}_n)$ we have:*

$$\sqrt{T_n - 1}[\eta_n^{T_n} - \eta_n](\varphi) \Rightarrow \mathcal{N}(0, \sigma_n^2(\varphi))$$

where, setting $\varphi_n = \varphi - \eta_n(\varphi)$

$$\begin{aligned} \sigma_n^2(\varphi) &= \eta_n(G_n) \eta_n(\varphi_n^2) + [\sigma_{n-1}^2(Q_n(\varphi_n))]/[\eta_n(G_n) \eta_{n-1}(G_{n-1})] \quad n \geq 2 \\ \sigma_1^2(\varphi) &= \eta_1(G_1) \eta_1((\varphi - \eta_1(\varphi))^2) \end{aligned}$$

or equivalently for any $n \geq 1$

$$\sigma_n^2(\varphi) = \eta_n(G_n) \sum_{q=1}^n \frac{\gamma_q(G_q)^2}{\gamma_n(G_n)^2} \eta_q([Q_{q,n}(\varphi_n) - \eta_q(Q_{q,n}(\varphi_n))]^2). \quad (4)$$

Proof. Our proof proceeds via induction. For the case $n = 1$, by Lemma B.1 we need only deal with the term

$$\sqrt{(N-1)\eta_1(G_1)}[\eta_n^{N-1} - \eta_1](\varphi).$$

Then one need only apply the CLT for i.i.d. bounded random variables; this yields the result with

$$\sigma_1^2(\varphi) = \eta_1(G_1) \eta_1((\varphi - \eta_1(\varphi))^2).$$

We assume the result for $n-1$ and consider n . Let $\varphi_n = \varphi - \eta_n(\varphi)$ then we have

$$\sqrt{T_n - 1}[\eta_n^{T_n} - \eta_n](\varphi) = \sqrt{T_n - 1}[\eta_n^{T_n} - \Phi_n(\eta_{n-1}^{T_{n-1}})](\varphi_n) + \sqrt{T_n - 1}\Phi_n(\eta_{n-1}^{T_{n-1}})(\varphi_n). \quad (5)$$

Now the first term on the R.H.S. of (5) can be written as

$$\begin{aligned} \sqrt{T_n - 1}[\eta_n^{T_n} - \Phi_n(\eta_{n-1}^{T_{n-1}})](\varphi_n) &= \sqrt{T_n - 1}[\eta_n^{T_n} - \Phi_n(\eta_{n-1}^{T_{n-1}})](\varphi_n) - \sqrt{(N-1)\eta_n(G_n)}[\eta_n^{N-1} - \Phi_n(\eta_{n-1}^{T_{n-1}})](\varphi_n) \\ &\quad + \sqrt{(N-1)\eta_n(G_n)}[\eta_n^{N-1} - \Phi_n(\eta_{n-1}^{T_{n-1}})](\varphi_n). \end{aligned} \quad (6)$$

In addition, the second term on the R.H.S. of (5) can be written as

$$\begin{aligned} \sqrt{T_n - 1}\Phi_n(\eta_{n-1}^{T_{n-1}})](\varphi_n) &= \left[\sqrt{\frac{T_n - 1}{T_{n-1} - 1}} - \sqrt{\frac{\eta_{n-1}(G_{n-1})}{\eta_n(G_n)}} \right] \sqrt{T_{n-1} - 1}\Phi_n(\eta_{n-1}^{T_{n-1}})](\varphi_n) \\ &\quad + \sqrt{T_{n-1} - 1} \sqrt{\frac{\eta_{n-1}(G_{n-1})}{\eta_n(G_n)}} \Phi_n(\eta_{n-1}^{T_{n-1}})](\varphi_n) \end{aligned} \quad (7)$$

By Lemma B.1 the first term on the R.H.S. of (6) converges in probability to zero. Also, by Lemma B.3, Theorem 3.1 (which provides a strong law of large numbers) and the induction hypothesis ($\eta_{n-1}(Q_n(\varphi_n)) = 0$), the first term on the R.H.S. of (7) converges in probability to zero. Thus, by a corollary to Slutsky's theorem, we can consider the weak convergence of

$$\sqrt{(N-1)\eta_n(G_n)}[\eta_n^{N-1} - \Phi_n(\eta_{n-1}^{T_{n-1}})](\varphi_n) + \sqrt{\frac{\eta_{n-1}(G_{n-1})}{\eta_n(G_n)}} \sqrt{T_{n-1} - 1}\Phi_n(\eta_{n-1}^{T_{n-1}})](\varphi_n) := A(N) + B(N).$$

We now consider the characteristic function:

$$\mathbb{E}[\exp\{it(A(N) + B(N))\}] = \mathbb{E}[\{\mathbb{E}[\exp\{itA(N)\}|\mathcal{F}_{n-1}] - e^{-\tilde{\sigma}_n^2(\varphi_n)t^2/2}\} \exp\{itB(N)\}] + e^{-\tilde{\sigma}_n^2(\varphi_n)t^2/2} \mathbb{E}[\exp\{itB(N)\}] \quad (8)$$

where

$$\tilde{\sigma}_n^2(\varphi_n) = \eta_n(G_n)\eta_n(\varphi_n^2)$$

and \mathcal{F}_{n-1} is the filtration generated by the particle system up-to time $n-1$. We deal with the limit of the expectations on the R.H.S. of (8) independently. We will show that $\mathbb{E}[\exp\{itA(N)\}|\mathcal{F}_{n-1}] - e^{-\tilde{\sigma}_n^2(\varphi_n)t^2/2}$ will converge in probability to zero, by using Theorem A3 of Douc & Moulines (2008). To that end, we note that for any $\varphi \in \mathcal{B}_b(\mathbb{E}_n)$, we have that

$$\frac{1}{N-1} \sum_{i=1}^{N-1} [\varphi(X_n^i) - \Phi(\eta_{n-1}^{T_{n-1}})(\varphi)]$$

will converge in probability to zero (for example, by controlling the second moment with the Marcinkiewicz-Zygmund (M-Z) inequality). Then by Theorem 3.1 as $\Phi(\eta_{n-1}^{T_{n-1}})(\varphi)$ converges almost surely to $\eta_n(\varphi)$ that $\eta_n^{N-1}(\varphi)$ will converge in probability to $\eta_n(\varphi)$. Using this result it follows easily that

$$\eta_n(G_n) \frac{1}{N-1} \sum_{i=1}^{N-1} [\varphi(X_n^i) - \Phi(\eta_{n-1}^{T_{n-1}})(\varphi)]^2$$

converges in probability to $\tilde{\sigma}_n^2(\varphi_n)$. This verifies the first condition of Theorem A3 of Douc & Moulines (2008) (eq. (31) of that paper). As φ is bounded, it is straightforward to verify the second (Lindeberg-type) condition of Theorem 13 of Douc & Moulines (2008) (eq. (32) of that paper). Thus, application of this latter theorem shows that $\mathbb{E}[\exp\{itA(N)\}|\mathcal{F}_{n-1}] - e^{-\tilde{\sigma}_n^2(\varphi_n)t^2/2}$ converges in probability to zero. Then by the induction hypothesis

$$B(N) \Rightarrow \mathcal{N}(0, [\sigma_{n-1}^2(Q_n(\varphi_n))]/[\eta_n(G_n)\eta_{n-1}(G_{n-1})]). \quad (9)$$

Thus

$$\{\mathbb{E}[\exp\{itA(N)\}|\mathcal{F}_{n-1}] - e^{-\tilde{\sigma}_n^2(\varphi_n)t^2/2}\} \exp\{itB(N)\} \rightarrow_{\mathbb{P}} 0.$$

Application of Theorem 25.12 of Billingsley (1995), shows that

$$\lim_{N \rightarrow \infty} \mathbb{E}[\{\mathbb{E}[\exp\{itA(N)\}|\mathcal{F}_{n-1}] - e^{-\tilde{\sigma}_n^2(\varphi_n)t^2/2}\} \exp\{itB(N)\}] = 0.$$

Thus, returning to (8), we consider $e^{-\tilde{\sigma}_n^2(\varphi_n)t^2/2} \mathbb{E}[\exp\{itB(N)\}]$. Noting (9) and again applying Theorem 25.12 of Billingsley (1995) we yield

$$\lim_{N \rightarrow \infty} \mathbb{E}[\exp\{itB(N)\}] = e^{-[\sigma_{n-1}^2(Q_n(\varphi_n))]/[\eta_n(G_n)\eta_{n-1}(G_{n-1})]t^2/2}.$$

Thus, we have proved that

$$\sqrt{T_n - 1}[\eta_n^{T_n} - \eta_n](\varphi) \Rightarrow \mathcal{N}(0, \sigma_n^2(\varphi))$$

where

$$\sigma_n^2(\varphi) = \eta_n(G_n)\eta_n(\varphi_n^2) + [\sigma_{n-1}^2(Q_n(\varphi_n))]/[\eta_n(G_n)\eta_{n-1}(G_{n-1})].$$

The verification of the formula (4) for the asymptotic variance follows standard calculations and is omitted. \square

Remark 3.1. *The formula for the asymptotic variance of the particle filter in Section 2.2, pp. 304 of Del Moral (2004) is*

$$\sum_{q=1}^n \frac{\gamma_q(1)^2}{\gamma_n(1)^2} \eta_q([Q_{q,n}(\varphi_n) - \eta_q(Q_{q,n}(\varphi_n))]^2). \quad (10)$$

Comparing to the asymptotic variance formula (4), this latter formula is certainly smaller if for each $1 \leq q < n$

$$\eta_q(G_q)^2 \leq \eta_n(G_n). \quad (11)$$

An alternative interpretation of (11) is if $\{V_n\}_{n \geq 1}$ is a Markov chain with transition kernels $\{M_n\}_{n \geq 1}$ then (11) is

$$\mathbb{P}(V_q \in B_q | v_1 \in B_1, \dots, v_{q-1} \in B_{q-1})^2 \leq \mathbb{P}(V_n \in B_n | v_1 \in B_1, \dots, v_{n-1} \in B_{n-1}).$$

Whilst this can be difficult to verify in general, if the spaces are $E_n = E$, $n \geq 1$, potentials $G_n = G$ for each $n \geq 1$ and the Markov kernels are such that for each $n \geq 1$, $x \in E$, $M_n(x, \cdot) = \nu(\cdot)$ for some $\nu \in \mathcal{P}(E)$, then the L.H.S. of (11) is $\nu(G)^2$ and the R.H.S. is $\nu(G)$; so in this ideal scenario, the new algorithm asymptotically outperforms the old one with regards to variance. In general, one might believe that (4) is smaller than the formula (10), as its leading term (when $q = n$) is smaller and the condition (11) can certainly hold in many examples.

Remark 3.2. *One can also prove a CLT for the estimate of the filter; a direct corollary is, under (M_1) , using Theorem 3.2, we have*

$$\sqrt{T_n - 1} \left[\frac{\eta_n^{T_n}(G_n \varphi)}{\eta_n^{T_n}(G_n)} - \frac{\eta_n(G_n \varphi)}{\eta_n(G_n)} \right] \Rightarrow \mathcal{N}(0, \tilde{\sigma}_n^2(\varphi))$$

where

$$\tilde{\sigma}_n^2(\varphi) = \sum_{q=1}^n \frac{\gamma_q(G_q)^2}{\gamma_n(G_n)^2 \eta_n(G_n)} \eta_q([Q_{q,n}(G_n[\varphi_n - \eta_n(G_n \varphi)]) - \eta_q(Q_{q,n}(G_n[\varphi_n - \eta_n(G_n \varphi)]))]^2).$$

In comparison, the asymptotic variance of the estimate in Theorem 4 of Le Gland & Oudjane (2006) (which differs to the one in this article) has asymptotic variance

$$\hat{\sigma}_n^2(\varphi) = \sum_{q=1}^n \frac{\gamma_q(1)\gamma_q(G_q)}{\gamma_n(G_n)^2} \eta_q([Q_{q,n}(G_n[\varphi_n - \eta_n(G_n \varphi)]) - \eta_q(Q_{q,n}(G_n[\varphi_n - \eta_n(G_n \varphi)]))]^2).$$

Thus there is an asymptotic difference between the two procedures. In general, our approach is better with regards to asymptotic variance if

$$\frac{\gamma_q(G_q)^2}{\gamma_n(G_n)^2 \eta_n(G_n)} \leq \frac{\gamma_q(1)\gamma_q(G_q)}{\gamma_n(G_n)^2}$$

or using the Markov chain interpretation in Remark 3.1:

$$\frac{\mathbb{P}(V_1 \in B_1, \dots, V_{q-1} \in B_{q-1}, V_q \in B_q)}{\mathbb{P}(V_n \in B_n | v_1 \in B_1, \dots, v_{n-1} \in B_{n-1})} \leq \mathbb{P}(V_1 \in B_1, \dots, V_{q-1} \in B_{q-1}).$$

In general, one cannot say which is preferable, but in the case: the spaces are $E_n = E$, $n \geq 1$, potentials $G_n = G$ for each $n \geq 1$ and the Markov kernels are such that for each $n \geq 1$, $x \in E$, $M_n(x, \cdot) = \nu(\cdot)$ for some $\nu \in \mathcal{P}(E)$, both the L.H.S. and R.H.S. of the inequality are equal.

3.4 Normalizing Constant

Define the estimate of the normalizing constant:

$$\gamma_n^{T_n}(\varphi) := \left\{ \prod_{p=1}^{n-1} \frac{N-1}{T_p-1} \right\} \eta_n^{T_n}(\varphi).$$

The technical results used in this Section can be found in Appendix C.

3.4.1 Unbiasedness

Proposition 3.1. *We have for any $n \geq 1$, $N \geq 2$ and $\varphi \in \mathcal{B}_b(\mathbf{E}_n)$, that*

$$\mathbb{E}[\gamma_n^{T_n}(\varphi)] = \gamma_n(\varphi).$$

Proof. The proof uses the standard Martingale difference decomposition in Del Moral (2004), with some additional expectation properties that need to be proved. The case $n = 1$ follows from Lemma C.1, so we assume $n \geq 2$. We remark that for $p \in \{2, \dots, n\}$:

$$\gamma_p^{T_p}(1)\Phi_p(\eta_{p-1}^{T_{p-1}})(Q_{p,n}(\varphi)) = \gamma_{p-1}^{T_{p-1}}(1)\eta_{p-1}^{T_{p-1}}(Q_{p-1,n}(\varphi))$$

and hence that

$$\gamma_n^{T_n}(\varphi) - \gamma_n(\varphi) = \sum_{p=1}^n \gamma_p^{T_p}(1)[\eta_p^{T_p} - \Phi_p(\eta_{p-1}^{T_{p-1}})](Q_{p,n}(\varphi)).$$

Then by Lemma C.1, it follows that

$$\mathbb{E}[\gamma_p^{T_p}(1)[\eta_p^{T_p} - \Phi_p(\eta_{p-1}^{T_{p-1}})](Q_{p,n}(\varphi)) | \mathcal{F}_{p-1}] = 0$$

and hence that

$$\mathbb{E}[\gamma_n^{T_n}(\varphi) - \gamma_n(\varphi)] = 0$$

from which we easily conclude the result. \square

3.4.2 Non-Asymptotic Variance Theorem

Below the term $\sum_{s=1}^n \frac{\hat{\delta}_s^{(m)} \hat{\beta}_s^{(m)}}{\eta_s(G_s)}$ is as in Cérou et al. (2011). The expressions and interpretations for $\hat{\delta}_s^{(m)} \hat{\beta}_s^{(m)}$ can be found in Section 3.1. In addition, $(\eta_n^{T_n})^{\odot 2}$ is the U -statistic that is formed from our empirical measure $\eta_n^{T_n}$ and $(\eta_n^{T_n})^{\otimes 2}$ is the corresponding V -statistic. In addition $(\gamma_n^{T_n})^{\otimes 2}(F) = \gamma_{n-1}^{T_{n-1}}(1)^2 (\eta_n^{T_n})^{\otimes 2}(F)$ for $F \in \mathcal{B}_b(\mathbf{E}^2)$.

Proposition 3.2. *Assume (\hat{H}_m) . Then for any $n \geq 2$, $N \geq 3$*

$$N > \sum_{s=1}^n \frac{\hat{\delta}_s^{(m)} \hat{\beta}_s^{(m)}}{\eta_s(G_s)} \Rightarrow \mathbb{E} \left[\left(\frac{\gamma_n^{T_n}(1)}{\gamma_n(1)} - 1 \right)^2 \right] \leq \frac{4}{N} \sum_{s=1}^n \frac{\tilde{\delta}_s \hat{\delta}_s^{(m)} \hat{\beta}_s^{(m)}}{\eta_s(G_s)}.$$

Proof. The result follows essentially from Cérou et al. (2011). To modify the proof to our set-up, we will prove that for $F : \mathbf{E}^2 \rightarrow \mathbb{R}_+$ (where the expectation on the L.H.S. is w.r.t. the stochastic process that generates the SMC algorithm)

$$\mathbb{E}[(\eta_n^{T_n})^{\otimes 2}(F)] \leq \left(\frac{N-1}{N-2} \right)^n \mathbb{E}_\xi [\eta_1^{\otimes 2} C_{\xi_1} Q_2^{\otimes 2} C_{\xi_2} \dots Q_n^{\otimes 2} C_{\xi_n}(F)] \quad (12)$$

where for each $n \geq 1$, independently

$$\mathbb{P}_\xi(\xi_n = 1) = 1 - \mathbb{P}_\xi(\xi_n = 0) = \frac{1}{N-1}$$

with corresponding joint expectation \mathbb{E}_ξ and $C_1(F)(x, y) = F(x, x)$, $C_0(F)(x, y) = F(x, y)$. Once (12) is proved this gives a verification of Lemma 3.2, eq. (3.3) of Cérou et al. (2011), given this, the rest of the argument then follows Proposition 3.4 of Cérou et al. (2011) and Theorem 5.1 and Corollary 5.2 in Cérou et al. (2011) (note that the fact that we have an upper-bound with $\alpha = 0$ (as in Cérou et al. (2011)) does not modify the result). We will write expectations w.r.t. the probability space associated to the particle system enlarged with the (independent) $\{\xi_n\}_{n \geq 1}$ as \mathbb{E}_ξ .

Thus, we consider the proof of (12). We have

$$\mathbb{E}[(\gamma_n^{T_n})^{\otimes 2}(F) | \mathcal{F}_{n-1}] = \gamma_n^{T_n}(1)^2 \mathbb{E}[(\eta_n^{T_n})^{\otimes 2}(F) | \mathcal{F}_{n-1}].$$

Now

$$\begin{aligned} \mathbb{E}[(\eta_n^{T_n})^{\otimes 2}(F) | \mathcal{F}_{n-1}] &= \mathbb{E} \left[\frac{T_n - 2}{T_n - 1} (\eta_n^{T_n})^{\odot 2}(F) + \frac{1}{T_n - 1} \eta_n^{T_n}(C(F)) \middle| \mathcal{F}_{n-1} \right] \\ &\leq \Phi_n(\eta_{n-1}^{T_{n-1}})^{\otimes 2}(F) + \frac{1}{N-1} \Phi_n(\eta_{n-1}^{T_{n-1}})(C(F)) \\ &\leq \left(\frac{N-1}{N-2} \right) \mathbb{E}_\xi [\Phi_n(\eta_{n-1}^{T_{n-1}})^{\otimes 2}(C_{\xi_n}(F)) | \mathcal{F}_{n-1}] \end{aligned}$$

where we have used $(T_n - 2)/(T_n - 1) \leq 1$, $1/(T_n - 1) \leq 1/(N - 1)$ and Lemmas C.2 and C.1 to obtain the second line. Thus we have that

$$\begin{aligned} \mathbb{E}[(\gamma_n^{T_n})^{\otimes 2}(F)|\mathcal{F}_{n-1}] &\leq \gamma_n^{T_n}(1)^2 \left(\frac{N-1}{N-2}\right) \mathbb{E}_\xi[\Phi_n(\eta_{n-1}^{T_{n-1}})^{\otimes 2}(C_{\xi_n}(F))|\mathcal{F}_{n-1}] \\ &\leq \gamma_{n-1}^{T_{n-1}}(1)^2 \left(\frac{N-1}{N-2}\right) \mathbb{E}_\xi[(\eta_{n-1}^{T_{n-1}})^{\otimes 2}(Q_n C_{\xi_n}(F))|\mathcal{F}_{n-1}]. \end{aligned}$$

Using the above inequality, one can repeat the argument inductively to deduce (12). This completes the proof of the Proposition. \square

Remark 3.3. *The significance of the result is simply that if*

$$\sup_s \frac{\hat{\delta}_s^{(m)} \hat{\beta}_s^{(m)}}{\eta_s(B_\epsilon(y_s))} < c$$

then if $N > cn$ the relative variance will be constant in n . This will be useful for the PMCMC algorithm in Section 4.2.

4 Numerical Implementation

4.1 ABC Filtering

4.1.1 Linear Gaussian model

To investigate the alive particle filter, we consider the following linear Gaussian state space model (with all quantities one-dimensional):

$$\begin{aligned} Z_n &= Z_{n-1} + V_n, \\ Y_n &= 2Z_n + W_n, \quad t \geq 1 \end{aligned}$$

where $V_n \sim \mathcal{N}(0, \sigma_v^2)$, and independently $W_n \sim \mathcal{N}(0, \sigma_w^2)$. Our objective is to fit an ABC approximation of this HMM; this is simply to investigate the algorithm constructed in this article.

4.1.2 Set up

Data are simulated from the (true) model for $T = 5000$ time steps and $\sigma_v^2 \in \{0.1, 1, 5\}$ and $\sigma_w^2 \in \{0.1, 1, 5\}$. For $n \in \{1, \dots, T\}$, if $p_n \geq \frac{1}{500}$, where $p_t \stackrel{\text{i.i.d.}}{\sim} \mathcal{U}_{[0,1]}$ (the uniform distribution on $[0, 1]$), we have $Y_n = c$, where $c \in \{80, 90, \dots, 140, 150\}$. Recall $B_\epsilon(y) = \{u : |u - y| < \epsilon\}$ and we consider a fixed sequence of ϵ which values belong to set $\{5, 10, 15\}$, i.e. $\epsilon \in \{5, 10, 15\}$. We compare the alive particle filter to the approach in Jasra et al. (2012).

The proposal dynamics are as described in Section 2.3. For the approach in Jasra et al. (2012), $N = 2000$ and we resample every time. For the alive particle filter, we used $N = 1500$ particles; this is to keep the computation time approximately equal. We also estimate the normalizing constant via the alive filter at each time step and compare it with ‘exact’ values obtained via the Kalman filter in the limiting case $\epsilon = 0$. To assess the performance in normalizing constant estimation, the relative variance is estimated via independent runs of the procedure.

Our results are constructed as two parts. In the first part, we compare the performance of two particle filters under different scenarios. In the second part, we focus on examples where the approach in Jasra et al. (2012) collapses. All results were averaged over 50 runs. We note that, with regards to the results in this Section and the approach in Le Gland & Oudjane (2004); generally similar conclusions can be drawn with regards to comparison to the approach in Jasra et al. (2012).

4.1.3 Part I

In this part, the analyses of the alive particle filter were completed in approximately 115 seconds and approximately 103 seconds were taken for the approach in Jasra et al. (2012) (which we just term the particle filter). Our results are shown in Figures 1-6.

Figure 1 displays the log relative error for the alive filter to the particle filter. We present the time evolution of the \mathbb{L}_1 log relative error between the ‘exact’ and estimated first moment. From our results, the mean log relative error for each panel is $\{0.06, 0.04, 0.07\}$. Figure 2 plots the absolute \mathbb{L}_1 error of the alive particle filter error across

time. These results indicate, in the scenarios under study, that both filters are performing about the same time with regards to estimating the filter. This is unsurprising as both methods use essentially the same information, and the outlying values do not lead to a collapse of the particle filter. In addition, the behaviour in Figure 2, which is predicted in Theorem 3.1 under strong assumptions, appears to hold in a situation where the state-space is non-compact.

In Figure 3, we show the time evolution of the log of the normalizing constant estimate for three approaches, i.e. Kalman filter (black ‘—’ line), new ABC filter (red ‘-.-’ line) and SMC method (blue ‘...’ line). Figure 4 displays the (log) relative variance of the estimate of the normalizing constant via the alive particle filter, when using the Kalman filter as the ground truth. In Figure 3, there is unsurprisingly a bias in estimation of the normalizing constant, as the ABC approximation is not exact, i.e. $\epsilon \neq 0$. In Figure 4 the linear decay in variance proven in Proposition 3.2 is demonstrated (although under a log transformation).

In Figure 5 and 6, we show the number of particles used at each time step (that is to achieve N alive particles) of the alive filter (Figure 5) and the number of alive particles for the standard particle filter (Figure 6). Both Figures illustrate the effect of outlying data, where the alive filter has to work ‘harder’ (i.e. assigns more computational effort), whereas the standard filter just loses particles.




Figure 1: Estimation error of the first moment for the linear state space model. Each panel displays (Log) the ratio of \mathbb{L}_1 error of the alive filter to old filter.




Figure 2: Estimation error of the first moment for the linear state space model using the alive particle filter (red ‘★’ indicates the x-axis position of outlier)

4.1.4 Part II

In this part, we keep the initial conditions the same as in the previous Section but change the value of ϵ . Instead of using $\epsilon \in \{5, 10, 15\}$, we set smaller values to ϵ , i.e. $\epsilon \in \{3, 6, 12\}$ (recall the smaller ϵ , the closer the ABC approximation is to the true HMM (Theorem 1 of Jasra et al. (2012))). This change makes the standard particle filter collapse whereas the alive filter does not have this problem. All results were averaged over 50 runs and our results are shown in Figures 7-8.




Figure 3: Estimated normalizing constant for the linear state space model: Kalman filter (black '—'), alive filter (red '---') and old filter (blue '...'). Each panel displays the estimated normalizing constant across time.




Figure 4: (Log) Relative variance of normalizing constant of alive filter to Kalman filter for the linear state space model. Each panel displays the relative variance across time.

In Figure 7, we present the true simulated hidden trajectory along with a plot of the estimated X_t given by the two particle filters across time when $(\sigma_v, \sigma_w) = (\sqrt{5}, \sqrt{5})$. As shown in Figure 7, the alive filter can provide better estimation versus the old particle filter. Figure 8 displays the log relative error of the alive filter to old particle filter, which supports the previous point made, with regards to estimation of the hidden state. Based upon the results displayed, the alive filter can provide good estimation results under the same conditions when the old particle filter collapses.

4.2 Particle MCMC

We now utilize the results in Propositions 3.1-3.2. In particular, Proposition 3.1 allows us to construct an MCMC method for performing static parameter inference in the context of ABC approximations of HMMs.

Recall Section 2.1. Our objective is to sample from the posterior density:

$$\pi(\theta|y_{1:n}) = \frac{\int_{\mathcal{X}^n} \prod_{k=1}^n g_\theta^\epsilon(y_k|z_k) f_\theta(z_k|z_{k-1}) dz_{1:n} \pi(\theta)}{\int_{\mathcal{X}^n \times \Theta} \prod_{k=1}^n g_\theta^\epsilon(y_k|z_k) f_\theta(z_k|z_{k-1}) dz_{1:n} \pi(\theta) d\theta} \quad (13)$$

where g_θ^ϵ , f_θ is as (1) and $\pi(\theta)$ is a prior probability density on Θ . Throughout the Section, we set $N \geq 2$, $\epsilon > 0$, but in general omit dependencies on these quantities. In practice, one often seeks to sample from an associated probability on the extended state-space $\mathbb{E}^n \times \Theta$

$$\tilde{\pi}(\theta, z_{1:n}, u_{1:n}|y_{1:n}) \propto \prod_{k=1}^n \mathbb{I}_{B_\epsilon(y_k)}(u_k) g_\theta(u_k|z_k) f_\theta(z_k|z_{k-1}) \pi(\theta).$$

It is then easily verified that for any fixed $\theta \in \Theta$

$$\pi(\theta|y_{1:n}) = \int_{\mathbb{E}^n} \tilde{\pi}(\theta, z_{1:n}, u_{1:n}|y_{1:n}) dz_{1:n} du_{1:n}.$$




Figure 5: Number of particles used for alive filter for the linear state space model. Each panel displays the number of particles across time.




Figure 6: Number of particles used for the particle filter of Jasra et al. (2012) for the linear state space model. Each panel displays the number of particles across time.

A typical way to sample from $\tilde{\pi}(\theta, z_{1:n}, u_{1:n} | y_{1:n})$ is via the Metropolis-Hastings method, with proposing to move from $(\theta, z_{1:n}, u_{1:n})$ to $(\theta', z'_{1:n}, u'_{1:n})$ via the probability density:

$$q(\theta' | \theta) \prod_{k=1}^n g_{\theta'}(u'_k | z'_k) f_{\theta'}(z'_k | z'_{k-1})$$

such a proposal removes the need to evaluate g_{θ} which is not available in this context. As is well known e.g. Andrieu et al. (2010), such procedures typically do not work very well and lead to slow mixing on the parameter space Θ . This proposal can be greatly improved by running a particle-filter (the particle marginal Metropolis-Hastings (PMMH) algorithm) as in Andrieu et al. (2010); that is a Metropolis-Hastings move that will first move θ , via $q(\theta' | \theta)$ and then run the algorithm in Section 2.2 picking a whole path, l , $x^l_{1:n} \in \mathbb{E}^n$ the sample used with a probability proportional to $G_n(x^l_n)$. Remarkably, this procedure yields samples from (13) via an auxiliary probability density; the details can be found in Andrieu et al. (2010), but the apparently *fundamental* property is that the estimate of the normalizing constant is unbiased. Note also that the sample from the Markov chain $(\theta, x^l_{1:n})$ also provides a sample from $\tilde{\pi}(\theta, z_{1:n}, u_{1:n} | y_{1:n})$.

As we have seen in the context of both theory and applications, it appears that the alive filter in Section 2.3 out-performs the standard one, for a given computational complexity. In addition, as seen in Proposition 3.1, the estimate of the normalizing constant is unbiased. It is therefore a reasonable conjecture that one can construct a new PMMH algorithm, with the alive particle filter investigated previously in this article and that this might perform better (in some sense) than the standard PMMH just described. We remark that the justification of this new PMMH follows from the statements in Andrieu & Vihola (2012) (see also Andrieu & Roberts (2009)) and Proposition 3.1, but we provide details for completeness.

4.2.1 New PMMH Kernel

We will define an appropriate target probability to produce samples from (13), but we first give the algorithm:




Figure 7: (a) ‘True’ Z_t and (b) estimated Z_t across time for the linear state space model, where red (‘-’) indicates the alive particle filter and black ‘.’ indicates the particle filter.




Figure 8: Estimation error of the first moment for the linear state space model. Each panel displays (Log) the ratio of L_1 error of alive filter to the particle filter.

1. Sample $\theta(0)$ from any absolutely continuous distribution. Then run the particle filter (with parameter value $\theta(0)$) in Section 2.3 up-to time n , storing $\gamma_{n+1}^N(1)$ (now denoted $\gamma_{n+1,\theta(0)}^N(1)$). Pick a trajectory $x_{1:n}^i(0)$, $i \in \{1, \dots, T_n(0) - 1\}$, with probability

$$\frac{G_n(x_n^i(0))}{\sum_{i=1}^{T_n(0)-1} G_n(x_n^i(0))}.$$

Set $i = 1$.

2. Propose $\theta'|\theta(i-1)$ from a proposal with positive density on Θ (write it $q(\theta'|\theta)$). Then run the particle filter (with parameter value θ') in Section 2.3 up-to time n , storing $\gamma_{n+1,\theta'}^N(1)$. Pick a trajectory $(x_{1:n}^i)'$ with probability

$$\frac{G_n((x_n^i)')}{\sum_{i=1}^{T_n(0)-1} G_n((x_n^i)')}.$$

Set $\theta(i) = \theta'$, $\gamma_{n+1,\theta(i)}^N(1) = \gamma_{n+1,\theta'}^N(1)$ with probability:

$$1 \wedge \frac{\gamma_{n+1,\theta'}^N(1)}{\gamma_{n+1,\theta(i-1)}^N(1)} \frac{\pi(\theta')q(\theta(i-1)|\theta')}{\pi(\theta(i-1))q(\theta'|\theta(i-1))}.$$

Otherwise set $\theta(i) = \theta(i-1)$, $\gamma_{n+1,\theta(i)}^N(1) = \gamma_{n+1,\theta(i-1)}^N(1)$, $i = i + 1$ and return to the start of 2.

For readers interested in the numerical implementation, they can skip to the next Section, noting that the θ samples will come from the posterior (13); this is now justified in the rest of the section.

We construct the following auxiliary target probability on the state-space:

$$\bar{\mathbb{E}} = \Theta \times \left(\bigcup_{T_1=N}^{\infty} \left(\mathbb{E}^{T_1} \times \{T_1\} \times \left(\bigcup_{T_2=N}^{\infty} \left(\mathbb{E}^{T_2} \times \{1, \dots, T_1-1\}^{T_2} \times \{T_2\} \times \dots \times \bigcup_{T_n=N}^{\infty} \left(\mathbb{E}^{T_n} \times \{1, \dots, T_{n-1}-1\}^{T_n} \times \{T_n\} \times \{1, \dots, T_n-1\} \right) \dots \right) \right) \right) \right).$$

Whilst the state-space looks complicated it corresponds to the static parameter and all the variables (the states and the resampled indices) sampled by the alive particle filter up-to time-step n and then just the picking of one of the final paths.

For $n \geq 2$ (we omit θ from our notation) define

$$\Psi_n \left(d(x_n^1, \dots, x_n^{T_n}), a_{n-1}^1, \dots, a_{n-1}^{T_n}, T_n | x_{n-1}^{1:T_{n-1}}, T_{n-1} \right) := \frac{\mathbb{I}_{S_n}(x_n^1, \dots, x_n^{T_n}, T_n) \binom{T_n-1}{N-1} \prod_{i=1}^{T_n} \frac{G_{n-1}(x_{n-1}^{a_{n-1}^i})}{\sum_{i=1}^{T_{n-1}-1} G_{n-1}(x_{n-1}^i)} M_n(x_{n-1}^{a_{n-1}^i}, dx_n^i)}{\sum_{T_n=N}^{\infty} \sum_{a_{n-1}^{1:T_n} \in \{1, \dots, T_{n-1}-1\}} \binom{T_n-1}{N-1} \int_{\mathbb{E}^{T_n}} \mathbb{I}_{S_n}(x_n^1, \dots, x_n^{T_n}, T_n) \left[\prod_{i=1}^{T_n} \frac{G_{n-1}(x_{n-1}^{a_{n-1}^i})}{\sum_{i=1}^{T_{n-1}-1} G_{n-1}(x_{n-1}^i)} M_n(x_{n-1}^{a_{n-1}^i}, dx_n^i) \right]}$$

where for $n \geq 1$

$$S_n = \{(u_n^1, \dots, u_n^{T_n}, T_n) \in \mathbb{Y}^n \times \{N, N+1, \dots\} : \sum_{i=1}^{T_n-1} \mathbb{I}_{B_\epsilon(y_n)}(u_n^i) = N-1 \cap u_n^{T_n} \in B_\epsilon(y_n)\}.$$

In addition, set

$$\Psi_1 \left(d(x_1^1, \dots, x_1^{T_1}), T_1 \right) := \frac{\mathbb{I}_{S_1}(x_1^1, \dots, x_1^{T_1}, T_1) \binom{T_1-1}{N-1} \prod_{i=1}^{T_1} M_n(x_0, dx_1^i)}{\sum_{T_1=N}^{\infty} \binom{T_1-1}{N-1} \int_{\mathbb{E}^{T_1}} \mathbb{I}_{S_1}(x_1^1, \dots, x_1^{T_1}, T_1) \left[\prod_{i=1}^{T_1} M_1(x_0, dx_1^i) \right]}.$$

Then the PMMH algorithm just defined samples from the target

$$\bar{\pi}(\theta, d(\mathbf{x}_1, \dots, \mathbf{x}_n), \mathbf{a}_{1:n-1}, l, T_{1:n} | y_{1:n}) \propto G_n(x_n^l) \gamma_{n+1, \theta}^N(1) \prod_{k=2}^n \Psi_k \left(d(x_k^1, \dots, x_k^{T_k}), a_{k-1}^1, \dots, a_{k-1}^{T_k}, T_k | x_{k-1}^{1:T_{k-1}}, T_{k-1} \right) \times \Psi_1 \left(d(x_1^1, \dots, x_1^{T_1}), T_1 \right) \pi(\theta).$$

where $\mathbf{a}_k = (a_k^1, \dots, a_k^{T_k})$, $\mathbf{x}_k = (x_k^1, \dots, x_k^{T_k})$ and $l \in \{1, \dots, T_n-1\}$. Using Proposition 3.1, one can easily verify that for any fixed $\theta \in \Theta$

$$\pi(\theta | y_{1:n}) = \int_{\bar{\mathbb{E}} \setminus \Theta} \bar{\pi}(\theta, d(\mathbf{x}_1, \dots, \mathbf{x}_n), \mathbf{a}_{1:n-1}, k, t_{1:n} | y_{1:n}).$$

Note also that the samples $(\theta, x_{1:n}^l)$ from $\bar{\pi}$ are marginally distributed according to $\tilde{\pi}(\theta, z_{1:n}, u_{1:n} | y_{1:n})$. The associated ergodicity of the new PMMH algorithm follows the construction in Andrieu et al. (2010) and we omit details for brevity.

4.2.2 Implementation on Real Data

We consider the following state-space model, for $n \geq 1$

$$\begin{aligned} Y_n &= \varepsilon_n \beta \exp(Z_n) \\ Z_n &= \phi Z_{n-1} + \sigma V_n \end{aligned}$$

where $\varepsilon_n \sim \mathcal{S}t(0, \xi_1, \xi_2, \xi_3)$ (a stable distribution with location parameter 0, scale ξ_1 , skewness parameter ξ_2 and stability parameter ξ_3) and $V_n \sim \mathcal{N}(0, c)$. We set $\theta = (\beta, c, \phi)$, with priors $c \sim \mathcal{IG}(2, 1/100)$, $\phi \sim \mathcal{IG}(2, 1/50)$ ($\mathcal{IG}(a, b)$ is an inverse Gamma distribution with mode $b/(a+1)$) and $\beta \sim \mathcal{N}(0, 10)$. Note that the inverse Gamma distributions have infinite variance.

We consider the daily (adjust closing) index of the S & P 500 index between 03/01/2011 – 14/02/2013 (533 data points). Our data are the log-returns, that is, if I_n is the index value at time n , $Y_n = \log(I_n/I_{n-1})$. The data

are displayed in Figure 9. The stable distribution may help us to more realistically capture heavy tails prevalent in financial data, than perhaps a standard Gaussian. In most scenarios, the probability density function of a stable distribution is intractable, which suggests that an ABC approximation might be a sensitive way to approximate the true model.




Figure 9: S & P 500 (a) index data and (b) (Log) Daily return

4.2.3 Algorithm setup

We consider two scenarios to compare the standard PMMH algorithm and the new one developed above. In the first situation we set $\xi_3 = 1.75$ and in the second, $\xi_3 = 1.2$, with $\xi_1 = \xi_2 = 1$ in both situations. In the first case, we make ϵ a suitable value as the data are not expected to jump off the same scale as the initial data. In the second, ϵ is significantly reduced; this is to illustrate a point about the algorithm we introduce. Both algorithms are run for about the same computational time, such that the new PMMH algorithm has 20000 iterations. The parameters are initialized with draws from the priors. The proposal on β is a normal random walk and for (c, ϕ) a gamma proposal centered at the current point with proposal variance scaled to obtain reasonable acceptance rates. We consider $N \in \{10, 100, 1000\}$ and for the new PMMH algorithm this value is lower to allow the same computational time.

4.2.4 Results

Our results are presented in Figures 10-13. In Figures 10-11 we can see the output in the case that $\xi_3 = 1.75$. For all cases, it appears that both algorithms perform very well; the acceptance rates were around 0.25 for each case. For the PMMH algorithm the average number of simulations of the data, per-iteration and data-point, were (1636, 745, 365) for $N \in \{1000, 100, 10\}$ respectively (recall we have modified N to make the computational time similar to the standard PMMH). For this scenario one would prefer the standard PMMH as the algorithmic performance is very good, with a removal of a random computation cost per iteration.

In Figures 12-13 the output when $\xi = 1.2$ is displayed. In Figure 12 we can see that the standard PMMH algorithm performs very badly, barely moving across the parameter space, whereas the new PMMH algorithm has very reasonable performance (Figure 13). In this case, ϵ is very small, and the standard SMC collapses very often, which leads to the undesirable performance displayed. We note that considerable effort was expended in trying to get the standard PMMH algorithm to work in this case, but we did not manage to do so (so we do not claim that the algorithm cannot be made to work). Note also that whilst these are just one run of the algorithms, we have seen this behaviour in many other cases and it is typical in these examples. The results here suggest that the new PMMH kernel might be preferred in difficult sampling scenarios, but in simple cases it does not seem to be required.

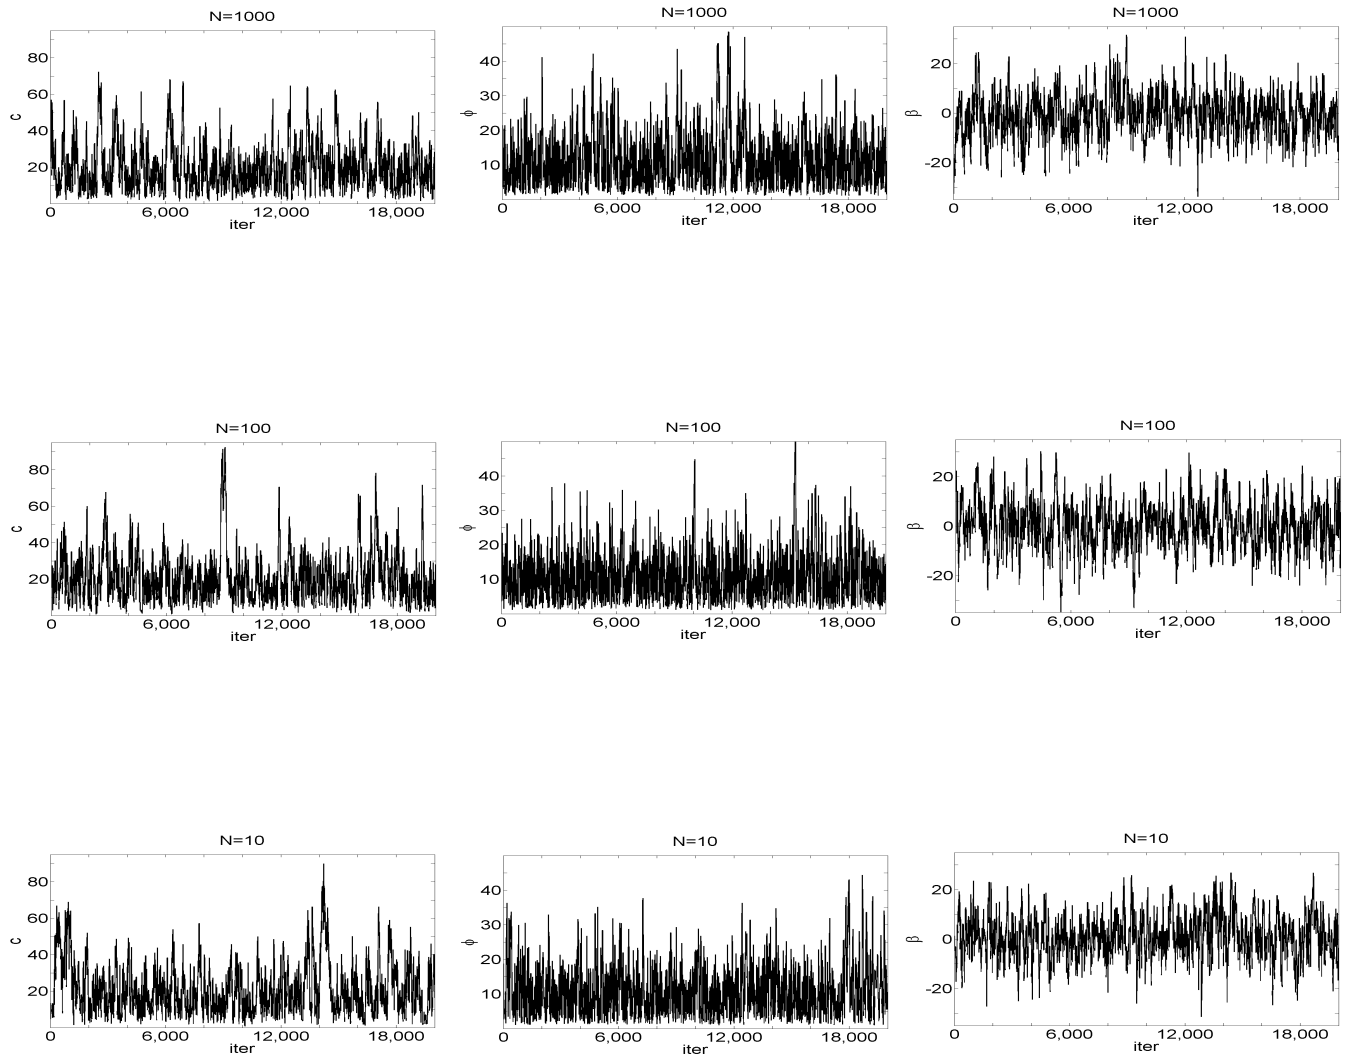



Figure 10: Trace plot of each parameter across iterations for a PMMH algorithm using the SMC algorithm in Section 2.2. Each row displays the samples with different N . Here $\xi_3 = 1.75$.



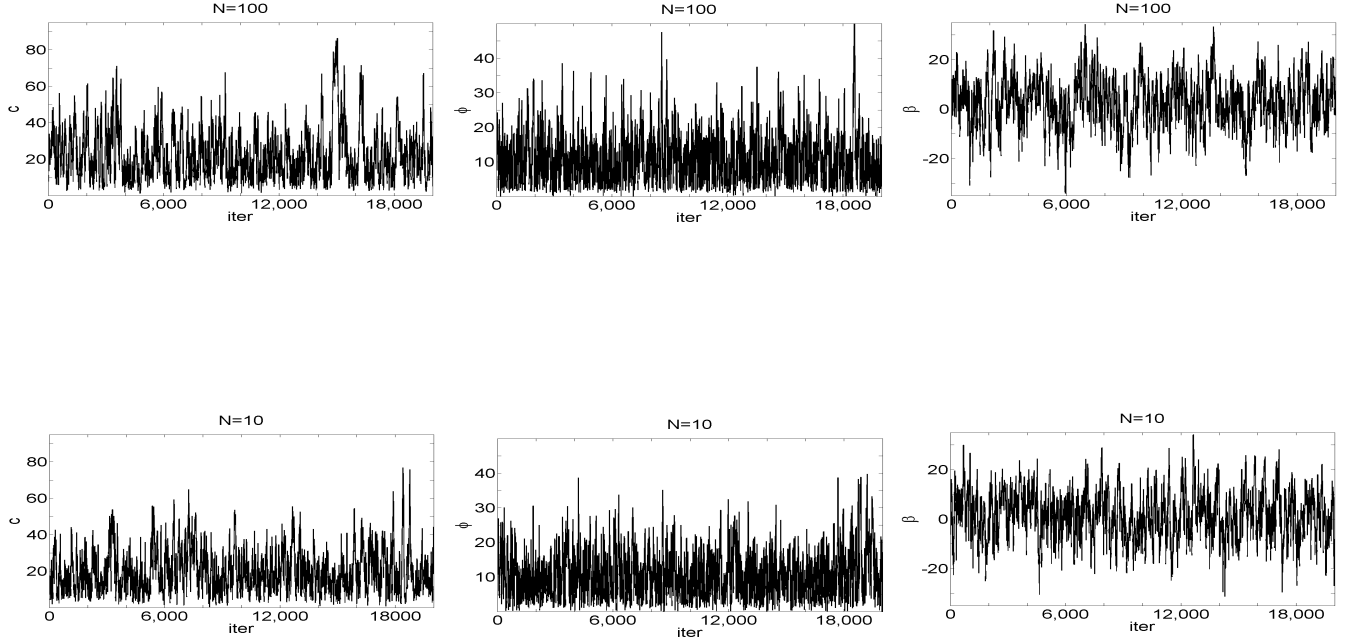
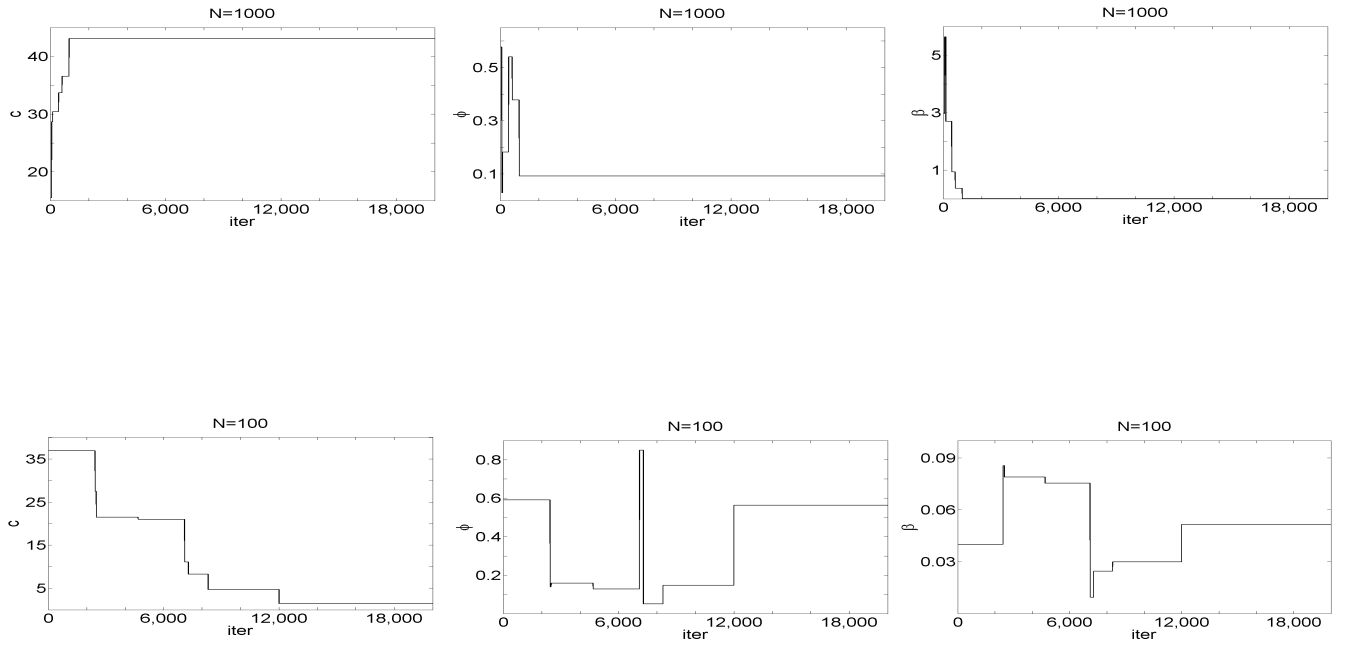


Figure 11: Trace plot of each parameter across iterations for a PMMH algorithm using the SMC algorithm in Section 2.3. Each row displays the samples with different N . Here $\xi_3 = 1.75$.






Figure 12: Trace plot of each parameter across iterations for a PMMH algorithm using the SMC algorithm in Section 2.2. Each row displays the samples with different N . Here $\xi_3 = 1.2$.

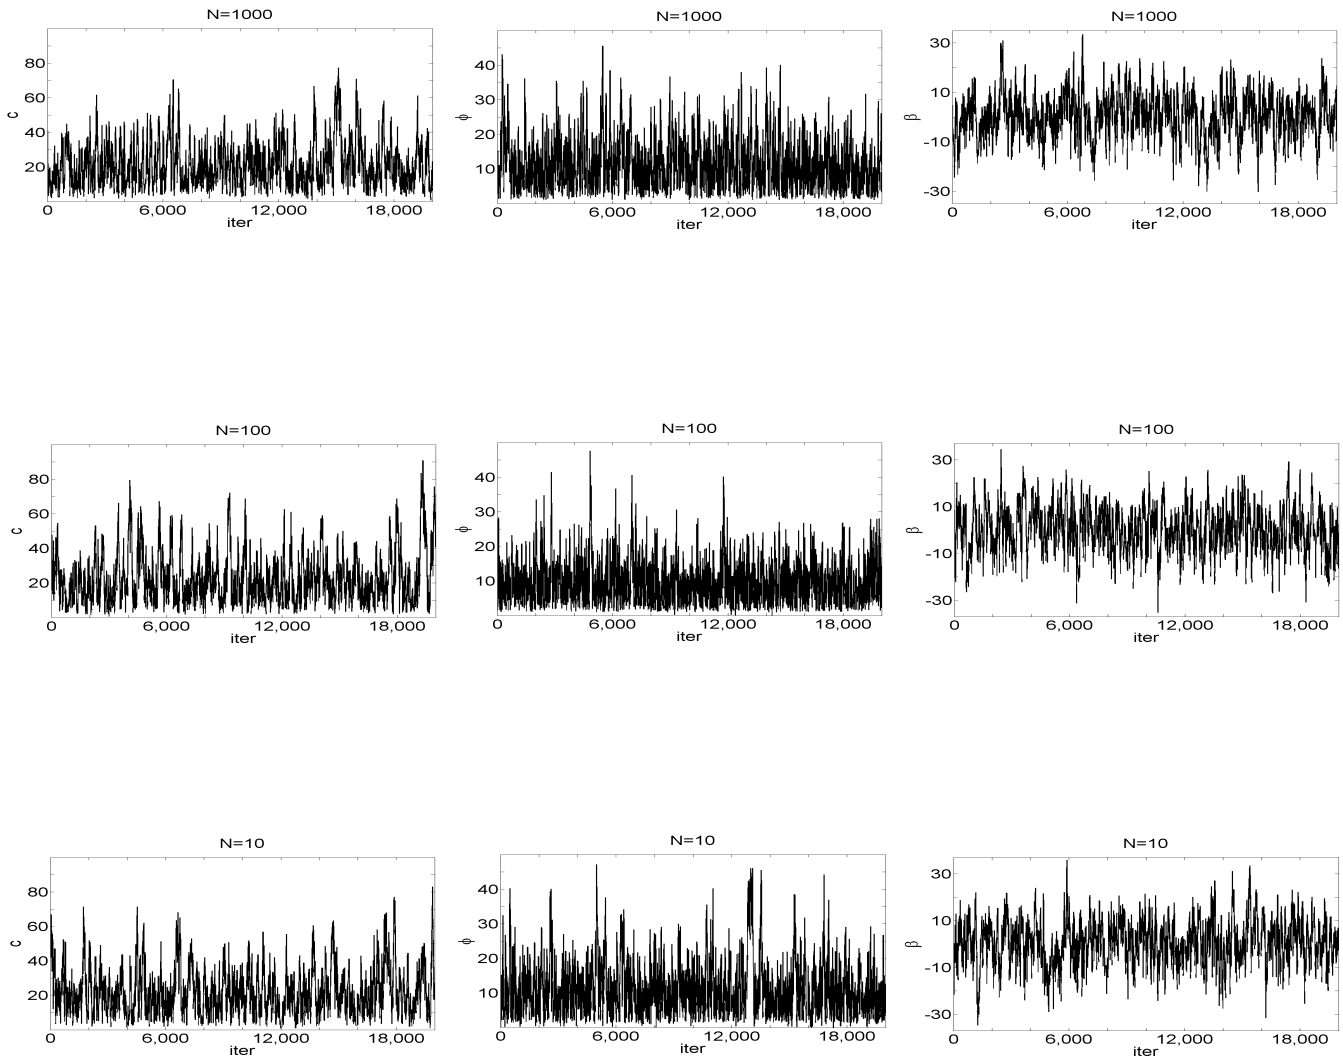


Figure 13: Trace plot of each parameter across iterations for a PMMH algorithm using the SMC algorithm in Section 2.3. Each row displays the samples with different N . Here $\xi_3 = 1.2$.

5 Summary

In this article we have investigated the alive particle filter; we developed and analyzed new particle estimates and derived new and principled MCMC algorithms. There are several extensions to the work in this article. Firstly, we have presented and analyzed the most standard particle filter; one can investigate more intricate filters commensurate with the current state of the art. Secondly, our theoretical results appear to hold under much weaker conditions than adopted (Section 4.1); one could extend the results in this direction. Thirdly, we have presented the most basic PMCMC algorithm; one can extend to particle Gibbs methods and beyond. Finally, one can also use the SMC theory in this article to interact with that of MCMC theory to investigate the performance of our PMCMC procedures.

Acknowledgements

The first author was supported by an MOE Singapore grant. We thank Gareth Peters for useful conversations on this work.

A Technical Results for the Predictor

The main result of this Section is below. Note we use the convention $\Phi_1(\eta_0^{T_0})(\varphi) = \eta_1(\varphi)$ and recall \mathcal{F}_n is the filtration generated by the particle system up-to time n .

Corollary A.1. *Assume (M_1) . Then for any $p \in [1, 4]$ there exists a $C_p < \infty$ such that for any $n \geq 1$, $N \geq 2$ and $\varphi \in \mathcal{B}_b(\mathbf{E}_n)$*

$$\mathbb{E}[|\eta_n^{T_n} - \Phi_n(\eta_{n-1}^{T_{n-1}})|(\varphi)|^p | \mathcal{F}_{n-1}]^{1/p} \leq \frac{C_p \|\varphi\|}{\sqrt{N-1}} \quad \mathbb{P} - a.s..$$

Proof. The case $n = 1$ follows directly from Lemma A.1 and (M_1) , so we consider $n \geq 2$. For $n \geq 2$, by Lemma C.1 $\mathbb{E}[\eta_n^{T_n}(\varphi) | \mathcal{F}_{n-1}] = \Phi_n(\eta_{n-1}^{T_{n-1}})(\varphi)$, so conditional upon \mathcal{F}_{n-1} we are in the setting of Lemma A.1. By (M_1) we can verify that $\Phi_n(\eta_{n-1}^{T_{n-1}})(\mathbf{B}_n) \wedge \Phi_n(\eta_{n-1}^{T_{n-1}})(\mathbf{B}_n^c) \geq c$ for some deterministic constant $1 > c > 0$ and hence application of Lemma A.1 proves the result. \square

A.1 Additional Technical Results

In the following section let $(\mathbf{F}, \mathcal{F})$ be a measurable space and X^1, X^2, \dots be i.i.d. random variables on \mathbf{F} associated to $\nu \in \mathcal{P}(\mathbf{F})$. Let $B \in \mathcal{F}$ be such that

$$\nu(B) \wedge \nu(B^c) \geq c$$

for some $1 > c > 0$. Let $N \geq 2$ and define

$$T := \inf\{n \geq 1 : \sum_{i=1}^n \mathbb{I}_B(X^i) \geq N\}.$$

Note that T is a negative Binomial random variable, with parameters N and success probability $\nu(B)$. We will consider some \mathbb{L}_p -properties of

$$\frac{1}{T-1} \sum_{i=1}^{T-1} \varphi(X^i)$$

with $\varphi \in \mathcal{B}_b(\mathbf{F})$. Expectations are written as \mathbb{E} . Note that one can follow the proof of Lemma C.1 to show that

$$\mathbb{E}\left[\frac{1}{T-1} \sum_{i=1}^{T-1} \varphi(X^i)\right] = \nu(\varphi).$$

We then have the following technical results which are used in the main text.

Lemma A.1. *For any $p \in [1, 4]$ there exist a $C_p < \infty$ such that for any, $N \geq 2$, and $\varphi \in \mathcal{B}_b(\mathbf{F})$*

$$\mathbb{E}\left[\left|\frac{1}{T-1} \sum_{i=1}^{T-1} [\varphi(X^i) - \nu(\varphi)]\right|^p\right]^{1/p} \leq \frac{C_p \|\varphi\|_\infty}{\sqrt{N-1}}.$$

Proof. Throughout C_p is a finite positive constant (that only depends upon p) whose value may change from line to line. WLOG we will assume that $\nu(\varphi) = 0$. By the Minkowski inequality

$$\begin{aligned} \mathbb{E} \left[\left| \frac{1}{T-1} \sum_{i=1}^{T-1} \varphi(X^i) \right|^p \right]^{1/p} &\leq \mathbb{E} \left[\left| \frac{1}{T-1} \sum_{i=1}^{T-1} \varphi(X^i) - \frac{N-1}{T-1} \frac{\nu(\mathbb{I}_B \varphi)}{\nu(B)} - \frac{T-N}{T-1} \frac{\nu(\mathbb{I}_{B^c} \varphi)}{\nu(B^c)} \right|^p \right]^{1/p} + \\ &\quad \mathbb{E} \left[\left| \frac{N-1}{T-1} \frac{\nu(\mathbb{I}_B \varphi)}{\nu(B)} + \frac{T-N}{T-1} \frac{\nu(\mathbb{I}_{B^c} \varphi)}{\nu(B^c)} \right|^p \right]^{1/p}. \end{aligned} \quad (14)$$

Lemma A.2 will control the second term on the R.H.S. so we focus on the first term on the R.H.S..

We have

$$\begin{aligned} &\mathbb{E} \left[\left| \frac{1}{T-1} \sum_{i=1}^{T-1} \varphi(X^i) - \frac{N-1}{T-1} \frac{\nu(\mathbb{I}_B \varphi)}{\nu(B)} + \frac{T-N}{T-1} \frac{\nu(\mathbb{I}_{B^c} \varphi)}{\nu(B^c)} \right|^p \right]^{1/p} = \\ &\mathbb{E} \left[\left| \frac{1}{T-1} \sum_{i=1}^{T-1} \mathbb{I}_B(X^i) \left[\varphi(X^i) - \frac{\nu(\mathbb{I}_B \varphi)}{\nu(B)} \right] + \frac{1}{T-1} \sum_{i=1}^{T-1} \mathbb{I}_{B^c}(X^i) \left[\varphi(X^i) - \frac{\nu(\mathbb{I}_{B^c} \varphi)}{\nu(B^c)} \right] \right|^p \right]^{1/p}. \end{aligned}$$

Now conditioning upon T (so that $N-1$ samples lie in B and $T-N$ lie in B^c and we subtract the conditional expectations of the (conditionally) independent random variables) and applying an appropriately modified version of the M-Z inequality (e.g. Chapter 7 of Del Moral (2004)) we have

$$\begin{aligned} &\mathbb{E} \left[\left| \frac{1}{T-1} \sum_{i=1}^{T-1} \varphi(X^i) - \frac{N-1}{T-1} \frac{\nu(\mathbb{I}_B \varphi)}{\nu(B)} + \frac{T-N}{T-1} \frac{\nu(\mathbb{I}_{B^c} \varphi)}{\nu(B^c)} \right|^p \right]^{1/p} \leq \\ &C_p \mathbb{E} \left[\frac{1}{(T-1)^{p/2+1}} \left((T-1)(\|\varphi\|_\infty(1 + \frac{1}{\nu(B)})^p + (T-N)(\|\varphi\|_\infty(1 + \frac{1}{\nu(B^c)})^p) \right) \right]^{1/p} \end{aligned}$$

where we are using the conditional distribution of X^1, \dots, X^{T-1} , given T . Setting $\bar{c} = (\|\varphi\|_\infty(1 + \frac{1}{\nu(B)})^p \vee (\|\varphi\|_\infty(1 + \frac{1}{\nu(B^c)})^p)$ we have that

$$\mathbb{E} \left[\left| \frac{1}{T-1} \sum_{i=1}^{T-1} \varphi(X^i) - \frac{N-1}{T-1} \frac{\nu(\mathbb{I}_B \varphi)}{\nu(B)} + \frac{T-N}{T-1} \frac{\nu(\mathbb{I}_{B^c} \varphi)}{\nu(B^c)} \right|^p \right]^{1/p} \leq C_p \bar{c}^{1/p} \mathbb{E} \left[\frac{1}{(T-1)^{p/2}} \left(1 + \frac{T-N}{T-1} \right) \right]^{1/p}$$

and then noting $1/(T-1) \leq 1/(N-1)$, $(1 + \frac{T-N}{T-1}) \leq 2$ and using $\nu(B) \wedge \nu(B^c) \geq c$ for dealing with \bar{c} we have proved that

$$\mathbb{E} \left[\left| \frac{1}{T-1} \sum_{i=1}^{T-1} \varphi(X^i) - \frac{N-1}{T-1} \frac{\nu(\mathbb{I}_B \varphi)}{\nu(B)} + \frac{T-N}{T-1} \frac{\nu(\mathbb{I}_{B^c} \varphi)}{\nu(B^c)} \right|^p \right]^{1/p} \leq \frac{C_p \|\varphi\|_\infty}{\sqrt{N-1}}.$$

Returning to (14) and using Lemma A.2, along with above result allows us to complete the proof. \square

Lemma A.2. For any $p \in [1, 4]$ there exist a $C_p < \infty$ such that for any, $N \geq 2$, and $\varphi \in \mathcal{B}_b(\mathbb{F})$

$$\mathbb{E} \left[\left| \frac{N-1}{T-1} \frac{\nu(\mathbb{I}_B \varphi)}{\nu(B)} + \frac{T-N}{T-1} \frac{\nu(\mathbb{I}_{B^c} \varphi)}{\nu(B^c)} - \nu(\varphi) \right|^p \right]^{1/p} \leq \frac{C_p \|\varphi\|_\infty}{\sqrt{N}}.$$

Proof. Throughout C_p is a finite positive constant (that only depends upon p) whose value may change from line to line. WLOG we will assume that $\nu(\varphi) = 0$, so that $\nu(\mathbb{I}_B \varphi) = -\nu(\mathbb{I}_{B^c} \varphi)$. Then we have that

$$\begin{aligned} \zeta(N, T) &:= \frac{N-1}{T-1} \frac{\nu(\mathbb{I}_B \varphi)}{\nu(B)} + \frac{T-N}{T-1} \frac{\nu(\mathbb{I}_{B^c} \varphi)}{\nu(B^c)} \\ &= \frac{\nu(\mathbb{I}_B \varphi)N}{\nu(B)(1-\nu(B))} \left[\frac{N-1+\nu(B)-T\nu(B)}{N(T-1)} \right]. \end{aligned}$$

Now, via Minkowski

$$\mathbb{E} [|\zeta(N, T)|^p]^{1/p} \leq \frac{|\nu(\mathbb{I}_B \varphi)|N}{\nu(B)(1-\nu(B))} \left\{ \mathbb{E} \left[\left| \frac{1}{(T-1)^p} \left| 1 - \frac{T\nu(B)}{N} \right|^p \right]^{1/p} + \mathbb{E} \left[\left| \frac{\nu(B)-1}{N(T-1)} \right|^p \right]^{1/p} \right\} \quad (15)$$

As $1/(T-1) \leq 1/(N-1)$, we will focus on controlling the term

$$\mathbb{E} \left[\left| 1 - \frac{T\nu(B)}{N} \right|^p \right]^{1/p}.$$

If Y^1, Y^2, \dots are independent $\mathcal{Geo}(\nu(B))$ random variables then

$$\mathbb{E} \left[\left| 1 - \frac{T\nu(B)}{N} \right|^p \right]^{1/p} = \nu(B) \mathbb{E} \left[\left| \frac{1}{N} \sum_{i=1}^N Y^i - \frac{1}{\nu(B)} \right|^p \right]^{1/p}$$

and applying an appropriately modified version of the M-Z inequality (e.g. Chapter 7 of Del Moral (2004)) we have

$$\mathbb{E} \left[\left| 1 - \frac{T\nu(B)}{N} \right|^p \right]^{1/p} \leq \frac{C_p}{\sqrt{N}} \left[\frac{(1 - \nu(B)(1 - \nu(B) + \nu(B)^2))}{\nu(B)^4} \right]^{1/p}$$

where we have used the fourth central moment of a Geometric random variable and C_p is a constant that only depends upon p (that is independent of ν or B). Returning to (15) and noting $\nu(B) \wedge \nu(B^c) \geq c$, we have shown that

$$\mathbb{E}[|\zeta(N, T)|^p]^{1/p} \leq C \|\varphi\|_\infty \frac{N}{N-1} \left[\frac{C_p}{\sqrt{N}} + \frac{1}{N} \right]$$

from which we can easily conclude. \square

B Technical Results for the CLT

Recall convergence in probability is written $\rightarrow_{\mathbb{P}}$ and N is going to ∞ . In addition, that the convention $\Phi_1(\eta_0^{T_0})(\varphi) = \eta_1(\varphi)$ is used and again recall \mathcal{F}_n is the filtration generated by the particle system up-to time n .

Lemma B.1. *Assume (M_1) . Then for any $n \geq 1$, $\varphi \in \mathcal{B}_b(\mathbb{E}_n)$ we have:*

$$\sqrt{T_n - 1}[\eta_n^{T_n} - \Phi_n(\eta_{n-1}^{T_{n-1}})](\varphi) - \sqrt{(N-1)\eta_n(G_n)}[\eta_n^{N-1} - \Phi_n(\eta_{n-1}^{T_{n-1}})](\varphi) \rightarrow_{\mathbb{P}} 0.$$

Proof. We give the proof for any $n \geq 2$; the case $n = 1$ follows a similar proof with only notational modifications. Throughout the proof $0 < C < \infty$ is a deterministic constant independent of n and N whose value may change from line to line. Our proof follows a similar construction to that found in pp. 369 of Billingsley (1995). To that end, we have

$$\begin{aligned} & \sqrt{T_n - 1}[\eta_n^{T_n} - \Phi_n(\eta_{n-1}^{T_{n-1}})](\varphi) - \sqrt{(N-1)\eta_n(G_n)}[\eta_n^{N-1} - \Phi_n(\eta_{n-1}^{T_{n-1}})](\varphi) = \\ & \sqrt{T_n - 1}[\eta_n^{T_n} - \Phi_n(\eta_{n-1}^{T_{n-1}})](\varphi) - (T_n - 1)\sqrt{\frac{\eta_n(G_n)}{(N-1)}}[\eta_n^{T_n} - \Phi_n(\eta_{n-1}^{T_{n-1}})](\varphi) + \end{aligned} \quad (16)$$

$$(T_n - 1)\sqrt{\frac{\eta_n(G_n)}{(N-1)}}[\eta_n^{T_n} - \Phi_n(\eta_{n-1}^{T_{n-1}})](\varphi) - \sqrt{(N-1)\eta_n(G_n)}[\eta_n^{N-1} - \Phi_n(\eta_{n-1}^{T_{n-1}})](\varphi). \quad (17)$$

In Lemma B.2, we have shown that (16) converges in probability to zero; hence, we focus upon (17).

To shorten the subsequent notations, we set

$$\begin{aligned} S_n^{T_n}(\varphi) &= (T_n - 1)[\eta_n^{T_n} - \Phi_n(\eta_{n-1}^{T_{n-1}})](\varphi) \\ S_n^{N-1}(\varphi) &= (N - 1)[\eta_n^{N-1} - \Phi_n(\eta_{n-1}^{T_{n-1}})](\varphi). \end{aligned}$$

Then let $1 > \varepsilon > 0$ be given, and consider:

$$\begin{aligned} & \mathbb{P} \left(\left| S_n^{T_n}(\varphi) - S_n^{N-1}(\varphi) \right| \geq \varepsilon \sqrt{\frac{N-1}{\eta_n(G_n)}} \right) \\ & \leq \mathbb{P} \left(\left| T_n - 1 - \frac{N-1}{\eta_n(G_n)} \right| \geq \varepsilon^3(N-1) \right) + \mathbb{P} \left(\max_{|k - \frac{N-1}{\eta_n(G_n)}| \leq \varepsilon^3(N-1)} \left| S_n^k(\varphi) - S_n^{N-1}(\varphi) \right| \geq \varepsilon(N-1)^{1/2} \right) \end{aligned}$$

$$\leq \mathbb{P}\left(\left|T_n - 1 - \frac{N-1}{\eta_n(G_n)}\right| \geq \varepsilon^3(N-1)\right) + 2\mathbb{P}\left(\max_{1 \leq k \leq \varepsilon^3(N-1)} |S_n^k(\varphi)| \geq \varepsilon(N-1)^{1/2}\right).$$

Now for the latter probability, one can condition upon \mathcal{F}_{n-1} and apply Kolmogorov's inequality noting that the conditional variance of $\varphi(x) - \Phi_n(\eta_{n-1}^{T_{n-1}})(\varphi)$ is deterministically upper-bounded by $\text{Osc}(\varphi)^2$. Hence, we have that

$$\begin{aligned} & \mathbb{P}\left(\left|S_n^{T_n}(\varphi) - S_n^{N-1}(\varphi)\right| \geq \varepsilon \sqrt{\frac{N-1}{\eta_n(G_n)}}\right) \\ & \leq \mathbb{P}\left(\left|T_n - 1 - \frac{N-1}{\eta_n(G_n)}\right| \geq \varepsilon^3(N-1)\right) + 2C\varepsilon. \end{aligned}$$

Noting Lemma B.3 and that we can make ε arbitrarily small, the proof is completed. \square

Lemma B.2. *Assume (M_1) . Then for any $n \geq 1$, $\varphi \in \mathcal{B}_b(\mathbf{E}_n)$ we have:*

$$\sqrt{T_n - 1}[\eta_n^{T_n} - \Phi_n(\eta_{n-1}^{T_{n-1}})](\varphi) - (T_n - 1)\sqrt{\frac{\eta_n(G_n)}{(N-1)}}[\eta_n^{T_n} - \Phi_n(\eta_{n-1}^{T_{n-1}})](\varphi) \rightarrow_{\mathbb{P}} 0.$$

Proof. We give the proof for any $n \geq 2$; the case $n = 1$ follows a similar proof with only notational modifications. Throughout the proof $0 < C < \infty$ is a deterministic constant independent of n and N whose value may change from line to line. We will prove that

$$\sqrt{T_n - 1}[\eta_n^{T_n} - \Phi_n(\eta_{n-1}^{T_{n-1}})](\varphi) - (T_n - 1)\sqrt{\frac{\eta_n(G_n)}{(N-1)}}[\eta_n^{T_n} - \Phi_n(\eta_{n-1}^{T_{n-1}})](\varphi)$$

will go to zero in \mathbb{L}_1 . To that end, we rewrite this expression as

$$A(N) := [\eta_n^{T_n} - \Phi_n(\eta_{n-1}^{T_{n-1}})](\varphi) \left[\frac{T_n - 1}{\sqrt{N-1}} \sqrt{\eta_n(G_n)} \right] \left[\sqrt{\frac{(N-1)}{(T_n - 1)\eta_n(G_n)}} - 1 \right].$$

To simplify the subsequent notations, we define

$$B(N) := \left[\sqrt{\frac{(N-1)}{(T_n - 1)\eta_n(G_n)}} - 1 \right].$$

Then a simple application of Hölder's inequality gives

$$\mathbb{E}[|A(N)|] \leq \mathbb{E}\left[\left|[\eta_n^{T_n} - \Phi_n(\eta_{n-1}^{T_{n-1}})](\varphi) \left[\frac{T_n - 1}{\sqrt{N-1}} \sqrt{\eta_n(G_n)} \right]\right|^{3/2}\right]^{2/3} \mathbb{E}[|B(N)|^3]^{1/3}.$$

We will show that:

1. $\mathbb{E}\left[\left|[\eta_n^{T_n} - \Phi_n(\eta_{n-1}^{T_{n-1}})](\varphi) \left[\frac{T_n - 1}{\sqrt{N-1}} \sqrt{\eta_n(G_n)} \right]\right|^{3/2}\right]^{2/3}$ is upper-bounded by a finite deterministic constant C that is independent of N .
2. $\lim_{N \rightarrow \infty} \mathbb{E}[|B(N)|^3]^{1/3} = 0$.

This will conclude the proof.

Proof of 1. We have, by another application of Hölder, that

$$\mathbb{E}\left[\left|[\eta_n^{T_n} - \Phi_n(\eta_{n-1}^{T_{n-1}})](\varphi) \left[\frac{T_n - 1}{\sqrt{N-1}} \sqrt{\eta_n(G_n)} \right]\right|^{3/2}\right]^{2/3} \leq \mathbb{E}[|\eta_n^{T_n} - \Phi_n(\eta_{n-1}^{T_{n-1}})](\varphi)|^3]^{1/3} \mathbb{E}\left[\left|\left[\frac{T_n - 1}{\sqrt{N-1}} \sqrt{\eta_n(G_n)} \right]^3\right|^{1/3}\right]$$

Application of Corollary A.1 gives

$$\mathbb{E}\left[\left|[\eta_n^{T_n} - \Phi_n(\eta_{n-1}^{T_{n-1}})](\varphi) \left[\frac{T_n - 1}{\sqrt{N-1}} \sqrt{\eta_n(G_n)} \right]\right|^{3/2}\right]^{2/3} \leq \frac{C_3 \|\varphi\|_{\infty}}{\sqrt{N-1}} \mathbb{E}\left[\left|\left[\frac{T_n - 1}{\sqrt{N-1}} \sqrt{\eta_n(G_n)} \right]^3\right|^{1/3}\right]. \quad (18)$$

Now turning to the expectation on the R.H.S. of the inequality, we have

$$\mathbb{E} \left[\left| \left[\frac{T_n - 1}{\sqrt{N - 1}} \sqrt{\eta_n(G_n)} \right]^3 \right| \right] = \frac{\eta_n(G_n)^{3/2}}{(N - 1)^{3/2}} \mathbb{E}[T_n^3 - 3T_n^2 + 3T_n - 1]$$

Using the fact that, via (M_1)

$$\Phi_n(\eta_n^{T_{n-1}})(\mathbf{B}_n) \wedge \Phi_n(\eta_n^{T_{n-1}})(\mathbf{B}_n^c) \geq c \quad (19)$$

for deterministic c and standard properties on raw moments of negative binomial random variables, it follows that

$$\mathbb{E}[T_n^3 - 3T_n^2 + 3T_n - 1] \leq CN^3.$$

Thus, we can show that

$$\mathbb{E} \left[\left| \left[\frac{T_n - 1}{\sqrt{N - 1}} \sqrt{\eta_n(G_n)} \right]^3 \right|^{1/3} \right] \leq C \eta_n(G_n)^{1/2} N^{1/2}.$$

Returning to (18), we have shown that

$$\mathbb{E} \left[\left| [\eta_n^{T_n} - \Phi_n(\eta_n^{T_{n-1}})](\varphi) \left[\frac{T_n - 1}{\sqrt{N - 1}} \sqrt{\eta_n(G_n)} \right]^{3/2} \right|^{2/3} \right] \leq C \eta_n(G_n)^{1/2}$$

which completes the proof of 1.

Proof of 2. By Lemma B.3 and the continuous mapping theorem, we have that $B(N) \rightarrow_{\mathbb{P}} 0$. Thus if we can show that for some $\delta > 0$, $\sup_{N \geq 1} \mathbb{E}[|B(N)|^{3(1+\delta)}] < +\infty$ this will allow us to conclude. For simplicity of calculation, we set $\delta = 1/3$. Then, using the fact that $1/(T_n - 1) \leq 1/(N - 1)$ we have

$$\mathbb{E}[|B(N)|^4] \leq \mathbb{E} \left[\left| 1 - \sqrt{\frac{\eta_n(G_n)(T_n - 1)}{N - 1}} \right|^4 \right].$$

On expanding the brackets and removing the negative terms, the expectation on the R.H.S. is upper-bounded by

$$1 + \frac{6\eta_n(G_n)}{N - 1} \mathbb{E}[T_n - 1] + \frac{\eta_n(G_n)^2}{(N - 1)^2} \mathbb{E}[T_n^2 - 2T_n + 1].$$

Using the conditional negative binomial property of T_n this expression is equal to

$$1 + \frac{6\eta_n(G_n)}{N - 1} \left\{ \mathbb{E} \left[\frac{N}{\Phi_n(\eta_n^{T_{n-1}})(G_n)} \right] - 1 \right\} + \frac{\eta_n(G_n)^2}{(N - 1)^2} \left\{ \mathbb{E} \left[\frac{N(1 - \Phi_n(\eta_n^{T_{n-1}})(G_n))}{\Phi_n(\eta_n^{T_{n-1}})(G_n)^2} + \frac{N^2}{\Phi_n(\eta_n^{T_{n-1}})(G_n)^2} - \frac{2N}{\Phi_n(\eta_n^{T_{n-1}})(G_n)} + 1 \right] \right\}.$$

Applying (19) we easily show that this latter expression is, uniformly in N , upper-bounded by a constant C . That is, we have shown that $\sup_{N \geq 1} \mathbb{E}[|B(N)|^{3(1+\delta)}] < +\infty$, which completes the proof of 2. This completes the proof. \square

Lemma B.3. Assume (M_1) . Then for any $n \geq 1$, we have:

$$\frac{T_n}{N} - \frac{1}{\eta_n(G_n)} \rightarrow_{\mathbb{P}} 0.$$

Proof. We give the proof for any $n \geq 2$; the case $n = 1$ follows a similar proof with only notational modifications. In Theorem 3.1, we have proved that $\Phi_n(\eta_n^{T_{n-1}})(\varphi)$ converges almost surely to $\eta_n(\varphi)$ for $\varphi \in \mathcal{B}_b(\mathbf{E})$. Thus we consider

$$\mathbb{E} \left[\left(\frac{T_n}{N} - \frac{1}{\Phi_n(\eta_n^{T_{n-1}})(G_n)} \right)^2 \right].$$

Now conditionally upon \mathcal{F}_{n-1} , T_n is a negative binomial random variable with success probability $\Phi_n(\eta_n^{T_{n-1}})(G_n)$, so writing $X_{i,n}(N)$ as (conditionally) independent geometric random variables with the same success probability, we have

$$\mathbb{E} \left[\left(\frac{T_n}{N} - \frac{1}{\Phi_n(\eta_n^{T_{n-1}})(G_n)} \right)^2 \right] = \mathbb{E} \left[\mathbb{E} \left[\left(\frac{1}{N} \sum_{i=1}^N [X_{i,n}(N) - \frac{1}{\Phi_n(\eta_n^{T_{n-1}})(G_n)}] \right)^2 \middle| \mathcal{F}_{n-1} \right] \right].$$

Applying the conditional version of the M-Z inequality on the R.H.S. of the inequality, we have the upper-bound:

$$\frac{C}{N} \mathbb{E} \left[\frac{1 - \Phi_n(\eta_{n-1}^{T_{n-1}})(G_n)}{\Phi_n(\eta_{n-1}^{T_{n-1}})(G_n)} \right]$$

recalling that $\Phi_n(\eta_{n-1}^{T_{n-1}})(G_n) \geq c$ for some deterministic constant c we conclude that

$$\mathbb{E} \left[\left(\frac{T_n}{N} - \frac{1}{\Phi_n(\eta_{n-1}^{T_{n-1}})(G_n)} \right)^2 \right] \leq \frac{C}{N}.$$

The proof is completed on recalling that $\Phi_n(\eta_{n-1}^{T_{n-1}})(G_n)$ converges almost surely to $\eta_n(G_n)$. \square

C Technical Results for the Normalizing Constant

Lemma C.1. *We have for any $n \geq 1$, $N \geq 2$ and $\varphi \in \mathbb{B}_b(\mathbb{E}_n)$, that*

$$\mathbb{E}[\eta_n^{T_n}(\varphi) | \mathcal{F}_{n-1}] = \Phi_n(\eta_{n-1}^{T_{n-1}})(\varphi)$$

where $\Phi_1(\eta_{-1}^{T_{-1}})(\varphi) = M_1(\varphi)$.

Proof. We have, for any $n \geq 1$, $N \geq 2$ that $T_n | \mathcal{F}_{n-1}$ is a Negative Binomial random variable with parameters $N - 1$ and success probability $\Phi_n(\eta_{n-1}^{T_{n-1}})(\mathbb{B}_n) = \Phi_n(\eta_{n-1}^{T_{n-1}})(G_n)$ and note that from Neuts & Zacks (1967) and Zacks (1980)

$$\mathbb{E} \left[\frac{N-1}{T_n-1} \middle| \mathcal{F}_{n-1} \right] = \Phi_n(\eta_{n-1}^{T_{n-1}})(\mathbb{B}_n). \quad (20)$$

Now,

$$\begin{aligned} \mathbb{E}[\eta_n^{T_n}(\varphi) | \mathcal{F}_{n-1}] &= \mathbb{E} \left[\frac{1}{T_n-1} \sum_{i=1}^{T_n-1} \varphi(X_n^i) \middle| \mathcal{F}_{n-1} \right] \\ &= \mathbb{E} \left[\left(\frac{1}{T_n-1} \right) \left\{ (N-1) \frac{\Phi_n(\eta_{n-1}^{T_{n-1}})(\varphi \mathbb{I}_{\mathbb{B}_n})}{\Phi_n(\eta_{n-1}^{T_{n-1}})(\mathbb{B}_n)} + (T_n - N) \frac{\Phi_n(\eta_{n-1}^{T_{n-1}})(\varphi \mathbb{I}_{\mathbb{B}_n^c})}{\Phi_n(\eta_{n-1}^{T_{n-1}})(\mathbb{B}_n^c)} \right\} \middle| \mathcal{F}_{n-1} \right] \\ &= \mathbb{E} \left[\frac{N-1}{T_n-1} \frac{\Phi_n(\eta_{n-1}^{T_{n-1}})(\varphi \mathbb{I}_{\mathbb{B}_n})}{\Phi_n(\eta_{n-1}^{T_{n-1}})(\mathbb{B}_n)} + \left(1 - \frac{N-1}{T_n-1} \right) \frac{\Phi_n(\eta_{n-1}^{T_{n-1}})(\varphi \mathbb{I}_{\mathbb{B}_n^c})}{\Phi_n(\eta_{n-1}^{T_{n-1}})(\mathbb{B}_n^c)} \right] \end{aligned}$$

where we have used the fact that there are $N - 1$ particles that are ‘alive’ and $T_n - N$ that will die and used the conditional distribution of the samples given T_n . Now by (20), it follows then that

$$\mathbb{E}[\eta_n^{T_n}(\varphi) | \mathcal{F}_{n-1}] = \Phi_n(\eta_{n-1}^{T_{n-1}})(\varphi \mathbb{I}_{\mathbb{B}_n}) + \Phi_n(\eta_{n-1}^{T_{n-1}})(\varphi \mathbb{I}_{\mathbb{B}_n^c}) = \Phi_n(\eta_{n-1}^{T_{n-1}})(\varphi)$$

which concludes the proof. \square

Lemma C.2. *We have for any $n \geq 2$, $N \geq 3$, $\varphi \in \mathbb{B}_b(\mathbb{E}_n^2)$:*

$$\mathbb{E}[(\eta_n^{T_n})^{\odot 2}(\varphi) | \mathcal{F}_{n-1}] = \Phi_n(\eta_{n-1}^{T_{n-1}})^{\otimes 2}(\varphi).$$

Proof. We have:

$$\begin{aligned} \mathbb{E}[(\eta_n^{T_n})^{\odot 2}(\varphi) | \mathcal{F}_{n-1}] &= \\ \mathbb{E} \left[\frac{(N-1)(N-2)}{(T_n-1)(T_n-2)} \middle| \mathcal{F}_{n-1} \right] \frac{\Phi_n(\eta_{n-1}^{T_{n-1}})^{\otimes 2}(\mathbb{I}_{\mathbb{B}_n^2} \varphi)}{\Phi_n(\eta_{n-1}^{T_{n-1}})^{\otimes 2}(\mathbb{B}_n^2)} + \end{aligned} \quad (21)$$

$$\mathbb{E} \left[2(N-1) \left(\frac{1}{T_n-1} - \frac{N-2}{(T_n-1)(T_n-2)} \right) \middle| \mathcal{F}_{n-1} \right] \frac{\Phi_n(\eta_{n-1}^{T_{n-1}})^{\otimes 2}(\mathbb{I}_{\mathbb{B}_n \times \mathbb{B}_n^c} \varphi)}{\Phi_n(\eta_{n-1}^{T_{n-1}})(\mathbb{B}_n) \Phi_n(\eta_{n-1}^{T_{n-1}})(\mathbb{B}_n^c)} + \quad (22)$$

$$\mathbb{E} \left[\left(1 - \frac{N-1}{T_n-1} - \frac{N-1}{T_n-2} + \frac{(N-1)^2}{(T_n-1)(T_n-2)} \right) \middle| \mathcal{F}_{n-1} \right] \frac{\Phi_n(\eta_{n-1}^{T_{n-1}})^{\otimes 2}(\mathbb{I}_{(\mathbb{B}_n^c)^2} \varphi)}{\Phi_n(\eta_{n-1}^{T_{n-1}})^{\otimes 2}((\mathbb{B}_n^c)^2)}. \quad (23)$$

The three terms on the R.H.S. arise due to the $(N-1)(N-2)$ different pairs of particles which land in \mathcal{B}_n^2 (21), the $2(N-1)(T_n-N)$ pairs of different particles which land in $\mathcal{B}_n \times \mathcal{B}_n^c$ (22) and the $(T_n-N)(T_n-N-1)$ different pairs of particles which land in $(\mathcal{B}_n^c)^2$ (23); the factors of $\Phi_n(\eta_{n-1}^{T_{n-1}})$ arise from the conditional distributions of the particles given T_n (recalling that conditional on \mathcal{F}_{n-1} , T_n is a negative binomial random variables parameters N and $\Phi_n(\eta_{n-1}^{T_{n-1}})(B_\epsilon(y_n))$).

Now for (21), we have from Neuts & Zacks (1967) and Zacks (1980) that

$$\mathbb{E}\left[\frac{(N-1)(N-2)}{(T_n-1)(T_n-2)} \middle| \mathcal{F}_{n-1}\right] = \Phi_n(\eta_{n-1}^{T_{n-1}})^{\otimes 2}(\mathcal{B}_n^2).$$

so that (21) becomes

$$\Phi_n(\eta_{n-1}^{T_{n-1}})^{\otimes 2}(\mathbb{I}_{\mathcal{B}_n^2} \varphi).$$

Recalling (20) and using the above result, (22) becomes

$$2\Phi_n(\eta_{n-1}^{T_{n-1}})^{\otimes 2}(\mathbb{I}_{\mathcal{B}_n \times \mathcal{B}_n^c} \varphi).$$

Finally, noting that for any $t \neq 1, 2$ $1/(t-2) = 1/(t-1) + 1/[(t-1)(t-2)]$, and thus using the above results that

$$\mathbb{E}\left[\frac{N-1}{T_n-2} \middle| \mathcal{F}_{n-1}\right] = \Phi_n(\eta_{n-1}^{T_{n-1}})^{\otimes 2}(\mathcal{B}_n) + \frac{\Phi_n(\eta_{n-1}^{T_{n-1}})^{\otimes 2}(\mathcal{B}_n)^2}{N-1}$$

it follows that (23) is equal to

$$\Phi_n(\eta_{n-1}^{T_{n-1}})^{\otimes 2}(\mathbb{I}_{(\mathcal{B}_n^c)^2} \varphi).$$

Hence we have shown

$$\begin{aligned} \mathbb{E}[(\eta_n^{T_n})^{\odot 2}(\varphi) | \mathcal{F}_{n-1}] &= \Phi_n(\eta_{n-1}^{T_{n-1}})^{\otimes 2}(\mathbb{I}_{\mathcal{B}_n^2} \varphi) + 2\Phi_n(\eta_{n-1}^{T_{n-1}})^{\otimes 2}(\mathbb{I}_{\mathcal{B}_n \times \mathcal{B}_n^c} \varphi) + \Phi_n(\eta_{n-1}^{T_{n-1}})^{\otimes 2}(\mathbb{I}_{(\mathcal{B}_n^c)^2} \varphi) \\ &= \Phi_n(\eta_{n-1}^{T_{n-1}})^{\otimes 2}(\varphi). \end{aligned}$$

□

References

- [1] ANDRIEU, C., DOUCET, A. & HOLENSTEIN, R. (2010). Particle Markov chain Monte Carlo methods (with discussion). *J. R. Statist. Soc. Ser. B*, **72**, 269–342.
- [2] ANDRIEU, A., ROBERTS, G. O. (2009). The pseudo-marginal approach for efficient Monte Carlo computations. *Ann. Statist.*, **37**, 697–725.
- [3] ANDRIEU, C. & VIHOLA, M. (2012). Convergence properties of pseudo-marginal Markov chain Monte Carlo algorithms. arXiv:1210.1484 [math.PR]
- [4] BILLINGSLEY, P. (1995). *Probability and Measure*. 2nd Edition. Wiley, New York.
- [5] CÉROU, F., DEL MORAL, P., FURON, T. & GUYADER, A. (2012). Sequential Monte Carlo for rare event estimation. *Statist. Comp.*, **22**, 795–808.
- [6] CÉROU, F., DEL MORAL, P. & GUYADER, A. (2011). A non-asymptotic variance theorem for un-normalized Feynman-Kac particle models. *Ann. Inst. Henri Poincaré*, **47**, 629–649.
- [7] DEAN, T. A., SINGH, S. S., JASRA, A. & PETERS G. W. (2010). Parameter estimation for Hidden Markov models with intractable likelihoods. Technical Report, University of Cambridge.
- [8] DEL MORAL, P. (2004). *Feynman-Kac Formulae*. Springer, New York.
- [9] DEL MORAL, P. & DOUCET, A. (2004). Particle motions in absorbing medium with hard and soft obstacles. *Stoch. Anal.*, **22**, 1175–1207.
- [10] DEL MORAL, P., DOUCET, A. & JASRA, A. (2012). An adaptive sequential Monte Carlo method for approximate Bayesian computation. *Statist. Comp.*, **22**, 1223–1237.

- [11] DOUC, R. & MOULINES, E. (2008). Limit theorems for weighted samples with applications to sequential Monte Carlo methods. *Ann. Statist.*, **36**, 2344– 2376.
- [12] JASRA, A., SINGH, S. S., MARTIN, J. S. & MCCOY, E. (2012). Filtering via approximate Bayesian computation. *Statist. Comp.*, **22**, 1223–1237.
- [13] LEE, A., ANDRIEU, C. & DOUCET, A. (2013). An active particle perspective of MCMC and its application to locally adaptive MCMC algorithms. Work in progress.
- [14] LE GLAND, F. & OUDJANE, N. (2004). Stability and uniform approximation of nonlinear filters using the Hilbert metric, and application to particle filters. *Ann. Appl. Probab.*, **14**, 144-187.
- [15] LE GLAND, F. & OUDJANE, N. (2006). A sequential particle algorithm that keeps the particle system alive. In *Stochastic Hybrid Systems : Theory and Safety Critical Applications*, (H. Blom & J. Lygeros, Eds), Lecture Notes in Control and Information Sciences **337**, 351–389, Springer: Berlin.
- [16] NEUTS, M. F. & ZACKS, S. (1967). On mixtures of χ^2 and F – distributions which yield distributions of the same family. *Ann. Inst. Stat. Math.*, **19**, 527–536.
- [17] ZACKS, S. (1980). On some inverse moments of negative-binomial distributions and their application in estimation. *J. Stat. Comp. & Sim.*, **10**, 163-165.